

Testing large scale cosmology with MeerKAT

Dissertation
zur
Erlangung des Doktorgrades (Dr. rer. nat.)
der
Mathematisch-Naturwissenschaftlichen Fakultät
der
Rheinischen Friedrich-Wilhelms-Universität Bonn

von
Jonah Daniël Wagenveld
aus
Zoeterwoude, Zuid-Holland, Niederlande

Bonn, 18.12.2023

Angefertigt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät der Rheinischen
Friedrich-Wilhelms-Universität Bonn

1. Gutachter: Prof. Dr. Michael Kramer
2. Gutachter: Prof. Dr. Frank Bertoldi
Tag der Promotion: 20.06.2024
Erscheinungsjahr: 2024

Through our eyes, the universe is perceiving itself. Through our ears, the universe is listening to its harmonies. We are the witnesses through which the universe becomes conscious of its glory, of its magnificence.

Alan Watts

Acknowledgements

When I first came to Bonn, my greatest fear was that it was going to be lonely, and that I had lost and would not be able to build up again a network of close friends. I can now happily say that never in my life have I been proven so wrong, to the point that I could not imagine being anywhere else and having such a good time. Needless to say, there are many people who have contributed to this, and though I will try to be complete here, it has been a long PhD, and some things will be lost in the depths of time.

I will start by appreciating my first set of friends and office mates, and lament the time that got stolen from us. Tímea, Shalini, and Lauren, thanks for being awesome office mates for that first half a year, I wish we could have had more time together. To the Pie gang, Charlie, Geetam, and Lauren, though we don't see each other often anymore, I know that I have made friends for life. Our Austria trip had so many memorable events, I will never forget it or stop talking about it. I'll be saving the Holunderschnaps for our future exploits. Though it only lasted a short while, Jaswanth, Shilpa, and Miquel, it was a hell of a thing to be shortly part of the self-proclaimed 'cool' office. Sorry for leaving you so early, but there was work to be done and time to get serious. Natasha, it was an absolute pleasure to be your office mate for half a year. Your love for statistics resonated well with mine, and I enjoyed our discussions a lot. To my eternal office mate Miquel, going all the way from working from home together to our current situation sharing an office, I could not wish for a better office mate and friend.

There are of course many others that have all added to my enjoyment at the institute. Mary, you are my example of good leadership and organisational skills, I have learned a lot from you. We had an amazing trip to Chile that I will never forget, and not just because of the food poisoning. Ferdinand, it has been great being your gym buddy, and it has always been a pleasure hosting your boardgaming sessions at my house. Arunima, your love for dutch croquettes truly warms my heart, as do your shouts and exclamations when I'm in the office. To Jaswanth, Surya, Joscha, Tim, Jompoj, Alexander, Andrei, Denisha, Marlon, Kathrin, Jędrzej, Kristen, Tasha, Eleni, Joey and the other wonderful people in our group, I say thank you for being so genuine and fun, something which has always made this research group special. This also goes for the amazing secretaries we have had over the years, Sylvia, Mariya, and Kira. Last but not least, to my close collaborator Shilpa, it has been an absolute pleasure working on our (secret) project, and I am very much looking forward to see where it leads us.

Of course, it would be hard to get where I am now without the advise and guidance from the people involved in my projects. I will especially thank Dominik and Sebastian for the great discussions on cosmology and the dipole, as well as Neeraj and the MALS collaboration for trusting in my abilities and allowing me to work with all of this amazing MeerKAT data. As I believe a research group reflects upon the quality of its leader, I would like to appreciate Michael for his excellent hiring choices (including me) that make our group what it is. I have been told not to thank a supervisor for doing

their job, as after all this is what they get paid to do. I will thus not thank Hans for supervising me. I will however thank him for all the things he does that go beyond the job description, especially the friendliness and lightheartedness he has displayed since the moment we first spoke on the phone when I was only a young master student. I will thank him for trusting me (sometimes more than I trust myself) and my instincts, allowing me to carve my own path through this project, and for hearing me when I had concerns or problems. Hans was one of the main advantages for me coming to Bonn, and I am happy to say I still feel that way.

Of course it is not only the wonderful people in Bonn who have been in my life for these past years. Especially during the early lockdown days, when I went back home, it was their presence and support that helped me during those times. To my high school friends, Mats, Luuk, Lauri, and Maud, I am glad that we are still friends after all these years, and I am especially thankful for the times you have visited Bonn and me feel a bit like I was home again. To my good friends Damiaan and Stephi, I hope that we will be able to see each other more in the future, but know that you are friends for life. To my Leiden University friends and hiking gang, Christiaan, Jelle, Auke, Joey, Couzy, Len, it is always great to see you guys again, and I especially enjoyed our amazing hiking trips together to Morocco and Montenegro. Of course, a special shoutout goes to Erik, it's always gonna be you and me baby, and even though we got separated during our PhD, it has felt like we were always in it together.

Finally, I will thank my family, my mom, my dad, and my sister. You have always had such unwavering faith in me while letting me walk my own path, and even though I sometimes set high standards for myself, I know that you will be proud of me no matter what I do. Thanks for always being there, for hosting me when I am back for holidays and back in the beginning of lockdown. And I would be remiss to leave out the newest addition to our family, my nephew Xavier. It is crazy to think that when I started my PhD you were nothing more than a clump of cells, and now you are a small person that I can talk to and play with. I hope that someday you will read this and be interested in what your uncle did back in the day.

Contents

1	Introduction	1
1.1	The landscape of modern cosmology	2
1.2	Observations of the cosmic radio dipole	4
1.3	Cosmological interpretations	6
1.4	The MeerKAT Absorption Line Survey	8
1.5	This thesis	10
2	The radio sky: from emission to source counts	11
2.1	On the nature and distribution of extragalactic radio sources	12
2.2	Radio observations and synthesis imaging	16
2.2.1	Observation and calibration	22
2.2.2	Imaging	24
2.2.3	Self-calibration	27
2.3	Image analysis	27
2.3.1	Primary beam correction	27
2.3.2	Source extraction and cataloguing	29
3	The most homogeneous of catalogues	31
3.1	Introduction	32
3.2	MALS data	33
3.2.1	Observations and calibration	33
3.2.2	Self-calibration and continuum imaging	35
3.2.3	Spectral index images	35
3.2.4	Primary beam correction	36
3.2.5	Assessment of calibration	38
3.2.6	Assessment of image quality	39
3.3	Source finding	41
3.3.1	Stokes I catalogues	42
3.3.2	Evaluation of individual pointings	45
3.4	Combined catalogue	50
3.4.1	Correcting residual primary beam effects	51
3.4.2	Cross-matching catalogues	52
3.4.3	General assessment of the complete catalogue	53
3.4.4	Resolved and unresolved sources	58
3.4.5	Catalogue columns	59

3.5	Source characteristics	60
3.5.1	Extended sources in PyBDSF	61
3.5.2	Spectral indices	62
3.5.3	Number counts	65
3.6	Towards the cosmic radio dipole	69
3.6.1	Compiling a homogeneous catalogue	70
3.6.2	Limitations and future prospects	71
3.7	Summary and conclusion	73
3.A	Cutouts of bright FRI and FRII sources	75
3.B	Table of sources	77
4	Rise of the Bayesian estimators	81
4.1	Introduction	82
4.2	Radio source counts and the dipole	83
4.3	Dipole estimators	84
4.3.1	Quadratic estimator	85
4.3.2	Poisson estimator	85
4.3.3	Poisson-rms estimator	86
4.3.4	Multi-Poisson estimator	87
4.3.5	Priors and injection values	87
4.4	Data	88
4.4.1	NVSS	88
4.4.2	RACS	90
4.4.3	Common masks and pixelation	91
4.4.4	Simulations	91
4.5	Results	92
4.5.1	Quadratic and Poisson estimators	93
4.5.2	Poisson-rms estimator	95
4.5.3	Combining RACS and NVSS	96
4.6	Discussion	98
4.6.1	The Poisson solution	98
4.6.2	Residual anisotropies in the data	99
4.6.3	Combining and splitting catalogues	102
4.6.4	Combining catalogues and cosmological considerations	104
4.7	Conclusion	105
5	Revenge of the systematics	107
5.1	MALS data	107
5.1.1	Noise properties	109
5.1.2	Source catalogues	111
5.2	A declination systematic	113
5.2.1	(u, v) -coverage effects	113
5.2.2	Fitting the data	117
5.3	Number counts	117

5.4	Preparing for a dipole measurement	120
5.4.1	MALS data	120
5.4.2	Estimators	121
5.4.3	Evaluating the estimators with simulations	122
5.5	MALS results	126
5.5.1	Declination estimator	126
5.6	Discussion	129
5.6.1	Declination systematics in other radio surveys	130
5.7	Conclusion	130
6	Outlook	133
6.1	Multi-wavelength synergies	133
6.2	Polarisation	134
6.3	Clustering	136
6.4	Isolating the kinematic dipole	137
7	Summary and conclusion	139
A	The radio number count dipole	143
A.1	Derivation	143
A.2	Advanced dipole likelihood	145
B	Paper I	149
C	Paper II	181
	Bibliography	197
	Summary	203
	Zusammenfassung	209
	Samenvatting	215

Introduction

The Copernican principle states that Earth, from which we are observing the Universe, does not occupy a special place in the Universe. It is named after Nicolaus Copernicus, who in 1543 published and popularised the heliocentric model of the Solar System, placing the Sun rather than the Earth at the centre of the Solar System (Copernicus, 1543). Removing Earth from the centre of the Solar System, and by extension the Universe, presented a paradigm shift that has informed how we view our place in the Universe ever since.

When developing his Theory of General Relativity, Albert Einstein extended his description to consider the relation between matter, energy, and space-time of the Universe at the largest scales (Einstein, 1917). Based on the Copernican principle, he used the 'Cosmological principle', which states that as a result of the laws of physics being the same everywhere in the Universe, we may assume that it is both homogeneous ('looks the same wherever you are') and isotropic ('looks the same in every direction') on the largest scales. He considered a simple Universe with uniformly distributed matter, and found that the solutions were unstable; the resulting Universe would either expand indefinitely or contract until it collapsed. This was an undesirable outcome, as Einstein intended this Universe to be static in time. His solution was to introduce a 'cosmological constant', Λ , intended to stabilise the system and thus create a static Universe. However, once evidence arose that the Universe might be expanding (Lemaître, 1927; Hubble and Humason, 1931), the cosmological constant was quickly discarded in favour of the original solution of Einstein's equations, which in fact produced this very result. Einstein's immediate reaction to try to 'fix' this apparent defect in his solutions reflected a broader sentiment in the scientific community. With cosmological theory and observations revealing the true scale and characteristics of the Universe, a new divide was created between competing models of the time evolution of the Universe. Taking the Copernican principle to the extreme, the 'perfect cosmological principle' states that we should not occupy a special place in time as well, and as such the Universe should be unchanging both in space and time. This model was called 'Steady state' cosmology, which posits that the Universe is forever unchanging on the largest scales (e.g. Hoyle, 1948; Bondi and Gold, 1948). This theory was directly at odds with the predictions of General Relativity, which implies evolution is a natural consequence of the interactions between matter and space-time. As observations at this time already suggested that the Universe was expanding, what we call now the Big Bang theory was the leading theory at the time, competing with the Steady state.

Although the final blow to Steady state theory was the discovery of the Cosmic Microwave Background (CMB, Penzias and Wilson, 1965), one of the major observational evidences against it

before came from a field that was only in its early developmental stage at the time: Radio Astronomy. It was after all only 20 years earlier that the first extraterrestrial radio emission, coming from the centre of the Milky Way, was discovered (Jansky, 1933). At this point in time, little was known of the sources that were seen at radio wavelengths, and because the telescopes at the time were not able to resolve most of them, they were dubbed ‘radio stars’. With a sample of 1906 sources, Ryle and Scheuer (1955) were able to show that they were isotropically distributed across the sky, making it highly likely that these sources were not contained within the Milky Way, but rather at cosmic distances. This interpretation made it possible to run an additional check. As the Steady state theory predicted uniformity in time as well as space, the number of sources seen further back in the Universe will not change, but will only on average appear fainter. As the number of sources increase with the cube of the lookback distance, and their brightness decreases with the square, the number of sources seen at lower brightness should increase as a power law with an index of $3/2$. However what Ryle and Scheuer (1955) saw was a much steeper increase of the number of sources at lower brightness. This had one obvious explanation, namely that in the past the density of these sources was larger than it is now, suggesting that the entire Universe was more dense as well. This was a strong argument for an expanding non-static Universe and in excellent agreement with the predictions of the Big Bang theory. In one of the first studies of the number counts of radio sources, the result had provided an important clue in setting the cosmological paradigm.

1.1 The landscape of modern cosmology

Today, the theory of General Relativity and its framework of describing the effects of gravity as an effect of the curvature of space-time has been proven time and again to correctly predict observed gravitational phenomena. Impressively, general relativity produces correct predictions about gravity just about everywhere we look, only breaking down on the smallest scales where quantum physics takes over. On the largest scales, it is well capable of describing the Universe at large. Built on the assumptions of the cosmological principle, the Friedmann-Lemaître-Robertson-Walker (FLRW) metric is an exact solution to Einstein’s field equations describing a homogeneous, isotropic, and non-static Universe, and is used as a basis for the most popular cosmological models at the current time. In this model, different constituents of the Universe, such as matter, radiation, and curvature, evolve differently with its expansion. The rate of expansion, the density, and the curvature of space-time are not predetermined in this model, and must be constrained by measurements of cosmological observables which are known to trace the structure and evolution of Universe. These observables are nowadays divided into ‘late time’ probes, constituting observations in the relatively nearby Universe, and ‘early time’, which describes the measurements performed using the CMB. With observations from both these probes, a model has been built up over the years based on the FLRW metric known as the lambda cold dark matter (Λ CDM) model. The name of this model already signifies the most important components of the Universe we live in. Cold dark matter describes matter which only seems to interact with matter and radiation through gravity, and makes up the majority of the matter in the Universe. Though the gravitational effects of dark matter can ostensibly be seen in many places and at many scales, a direct detection of its constituent remain elusive. Λ is the same cosmological constant once introduced by Einstein, now reintroduced into the model to account for the fact that the expansion of the Universe seems to be accelerating at late times, as was first found by observations of type Ia supernovae (Riess et al., 1998; Perlmutter et al., 1999). So far, it is not known what causes




Figure 1.1: Temperature map of the cosmic microwave background at 100 GHz in galactic coordinates. With the monopole subtracted but no other corrections made, the galactic plane and the dipole are the most prominent features visible.

the acceleration of the expansion of the Universe, so cosmologists refer to this rather mysterious force as ‘Dark energy’. An addition to most cosmological models is inflation, which takes place shortly (between 10^{-36} and 10^{-33} seconds) after the Big Bang and causes the Universe to rapidly and exponentially expand. Inflation helps explain why the observed Universe, and especially the CMB, appears homogeneous and isotropic even though parts of it have not been in causal contact with each other within the lifetime of the Universe. It also offers a satisfactory explanation of why the space-time curvature of the Universe appears to be almost flat. However, the mechanism behind inflation is not known.

While late time and early time measurements largely agreed on this state of the Universe, some remarkable differences have also been found. Of these ‘cosmological tensions’ the most well known one is the Hubble tension. It describes the increasingly significant difference between the value of the Hubble constant, H_0 , which describes the current rate of expansion of the Universe, inferred from early time experiments, and its value measured by late time observations. With increasing quality of data and processing we are now reaching the era of high precision cosmology, and in recent years this tension has critically reached a 5σ significance, indicating that the discrepancy is extremely unlikely to be a statistical fluke. A similar trend is seen across a wide range of late time cosmological probes, with such precise estimates that they become irreconcilable with predictions from early time cosmology.

It is one of these discrepancies that will be the focus of this thesis, and though it is a lesser known one, it presents an intriguing puzzle of measurements that seem impossible to coexist. In fact, it might

even call into question the validity of the cosmological principle, which sits at the very foundation of our current cosmological models. The CMB, an exceptional example of the large scale isotropy and homogeneity which we expect to see in the Universe, reveals much through the small density perturbations that can be seen at all scales. At first sight however, these are all hidden by a set of large anisotropies that have to be accounted for first. Figure 1.1 shows the two prominent anisotropies at CMB wavelengths as emission from the Galactic plane and a dipole. We will focus on the latter, which is asserted to arise quite simply from the movement of the observer with respect to the cosmic background. This can be understood as the observers' movement causing a Doppler effect on the CMB, blueshifting and thus seemingly increasing the temperature in the direction of motion, and causing the opposite in the other direction. With only the CMB measurement, this seems like an open and shut case with no further interesting features.

When we try to make a similar measurement at different wavelengths and tracers however, some unexpected results are obtained. With most bright radio sources being at cosmologically significant distances, they serve as alternative background frame from the CMB, expected to be completely isotropic. The movement of the observer induces several relativistic effects which has a net effect of creating a dipole in the number counts of these radio sources. From the CMB dipole, we can derive a solar system velocity of $v = 369.82 \pm 0.11$ km/s (Aghanim et al., 2020), so with any measurement of the radio dipole we would expect to find a similar value. However, it has been found in recent years that radio source counts imply a velocity that is at least twice as large. Taking these measurements at face value, they obviously cannot be true at the same time, which implies that either radio sources have a different reference frame than the CMB, or something other than our velocity might be adding to the dipole. Either way, a large scale anisotropy must be present to induce this, which is at odds with the cosmological principle and consequently the foundations upon which cosmology is built. It's safe to say that in order to make such a claim, we must be absolutely certain that our radio observations are free of biases and systematic effects. This is not always clear, as many have done this measurement using available catalogues of radio sources, which have not been created with this specific goal in mind.

1.2 Observations of the cosmic radio dipole

A measurement of Earth's velocity with respect to the cosmic background using the number counts of radio sources was first proposed by Ellis and Baldwin (1984). Due to the effects of relativistic aberration (Einstein, 1905), Doppler boost, and Doppler shift (Doppler, 1842), the movement of the observer will cause a dipole in the observed number density of radio sources. The dipole created due to these effects is also called the kinematic dipole, and as observed in the number density of radio sources, has an amplitude of

$$\mathcal{D} = [2 + x(1 + \alpha)]\beta. \quad (1.1)$$

In this equation, x is the power law index of the flux density distribution of sources and α represents the spectral index of sources¹. Finally, $\beta = v/c$ represents the velocity of the observer as a fraction of the speed of light. Assuming a velocity equal to that measured from the CMB, $v = 370$ km/s, the number of radio sources needed for a significant (3σ) measurement of the dipole would be $n \approx 2 \times 10^5$, adequately distributed along the axis of the dipole. At that time, no survey had the required sensitivity

¹ For a full derivation of this formula, see Section A.1

and sky coverage to reach this number. Crawford (2009) would recompute the number of required sources with a more advanced approach, concluding that $n \approx 2 \times 10^6$ would be necessary for a significant measurement of the dipole. As luck would have it, the Ellis and Baldwin (1984) estimate would be closer to the true value, by virtue of the radio dipole having a much higher amplitude than expected.

The first survey with which a successful significant measurement of the radio dipole was the National Radio Astronomy Observatory (NRAO) Very Large Array (VLA) Sky Survey (NVSS, Condon et al., 1998). The NVSS covers 82% of the full sky, including the entire northern hemisphere and the southern hemisphere down to a declination of -40° . The full catalogue contains 1.8 million sources, however several stringent cuts in the catalogue have to be made before it suitable for a dipole measurement. As sources inside the Milky Way are of no use to a dipole measurement, the entire Galactic plane needed to be removed. Furthermore, several systematic effects had an impact on the source density in different parts of the survey, which will introduce a bias in a dipole measurement if unaccounted for. The origin of such effects can differ per survey, and can be difficult to trace given the many processing steps which are performed between the collection of radio data and a catalogue of sources. In the case of NVSS, such an effect is introduced by the usage of different VLA array configurations to observe different parts of the sky. This causes an artificial difference in source density at different declinations, as was noted in the first measurement of the cosmic radio dipole with the NVSS by Blake and Wall (2002). These and other systematic effects mostly affect the faintest sources in the survey, making it necessary to cut out these sources, consequently reducing the number of sources that can be used for a dipole estimate. Therefore, Blake and Wall (2002) cut all sources in the Galactic plane and known nearby galaxies, and measured the dipole employing several flux density cuts, with the lowest cut at 10 mJy. With this, the number of sources used for the dipole measurement was reduced to around 4×10^5 , still enough to perform a measurement. Albeit it with considerably large uncertainties, the measured dipole was consistent, both in terms of amplitude and direction, with the CMB expectation.

Nearly ten years later, Singal (2011) revisited the NVSS and found that the radio number count dipole, though consistent in terms of direction with the CMB, exhibited an amplitude four times larger than what would be expected given the CMB dipole. This triggered a great number of other measurements using several large sky surveys, including the Westerbork Northern Sky Survey (WENSS, Rengelink et al., 1997), the Tata Institute for Fundamental Research (TIFR) Giant Metrewave Radio Telescope (GMRT) Sky Surveys first alternative data release (TGSS ADR1, Intema et al., 2017), and the Sydney University Molonglo Sky Survey (SUMSS, Mauch et al., 2003). Though some variation was observed depending on employed estimator, survey or masking strategy, the radio dipole was found to be generally consistent in terms of direction, but two to six times larger in terms of amplitude than the CMB dipole (e.g. Rubart and Schwarz, 2013; Tiwari et al., 2015; Colin et al., 2017; Siewert et al., 2021). However, lack of sources and systematic effects in surveys other than NVSS hampered the perceived reliability of these measurements. This is especially clear in measurements of the dipole with TGSS ADR1, which yields dipole amplitudes ten times higher than the CMB dipole amplitude, which is at least three times higher than any dipole amplitude measured from NVSS (Bengaly et al., 2018; Singal, 2019; Siewert et al., 2021). Siewert et al. (2021) suggests this might signal a frequency dependence of the dipole amplitude, however others attribute it to systematics originating from the flux scale of TGSS (Tiwari et al., 2019; Secrest et al., 2022). Darling (2022) combines the VLA Sky Survey (VLASS, Lacy et al., 2020) and the Rapid Australian Square Kilometre Array Pathfinder (ASKAP) Continuum Survey (RACS, McConnell et al., 2020) for a dipole measurement, surprisingly finding a dipole which is consistent both in terms of amplitude and direction with the CMB. This

result is however called into question in Secrest et al. (2022), who argue that combining catalogues in such a way is problematic, because it ignores the differences between catalogues introduced by any number of systematic effects, depending on frequency, angular resolution, and calibration to name a few. Furthermore, dipole measurements on the individual catalogues should yield results consistent with these results, but these are not performed in Darling (2022). Singal (2023) does perform dipole measurements on the individual catalogues finding that, in line with earlier results, the dipole amplitudes are once again higher than the CMB dipole.

Most recently, measurements of the dipole have been extended to other wavelengths and observables. Using 1.36 million quasars observed by the Wide-field Infrared Survey Explorer (WISE, Wright et al., 2010), Secrest et al. (2021) measure the dipole on a sample that is almost completely independent of previous measurements with radio sources, and find an amplitude twice as large as the CMB expectation, with a significance of 4.9σ . This measurement is repeated and combined with a NVSS measurement in Secrest et al. (2022), increasing the significance to 5.2σ . A Bayesian analysis of the sample is performed in Dam et al. (2023), achieving a statistical significance of 5.7σ . All these measurements find agreement between the number count dipole and the CMB dipole in terms of direction, but reject the CMB dipole with high statistical significance in terms of amplitude. As Secrest et al. (2022) remove sources from the WISE sample that have an NVSS counterpart, the infrared quasars are an independent sample from the radio sources. Furthermore, given that radio surveys are carried out with ground-based observatories and WISE is a space-based observatory, the systematics affecting these samples are expected to be completely different. As such, the excess dipole amplitude originating from a statistical fluke or as resulting from systematic effects is extremely unlikely at this point.

1.3 Cosmological interpretations

Though plenty of measurements of the dipole have already been performed, a satisfying explanation of the excess dipole signal remains elusive. The excess amplitude of the number count dipole might not be kinetic at all, but could even be a true anisotropy in the distribution of the measured sources. The dipole observed in number counts has two possible contributions

$$\mathcal{D} = \mathcal{D}_{kin} + \mathcal{D}_{int}, \quad (1.2)$$

where \mathcal{D}_{kin} represents the kinematic dipole from Equation 1.1, and \mathcal{D}_{int} an intrinsic dipole anisotropy in the distribution of the observed source population. Though it is possible to measure the kinetic and intrinsic contributions to the dipole separately (Tiwari et al., 2015; Nadolny et al., 2021), it requires much more data than is currently available. At the moment, our best bet is to look at other strange measurements which might signal towards a cause common with that of the excess number count dipole.

In the standard Λ CDM model, isotropy and homogeneity is required on the largest scales, both in terms of matter density and in terms of expected velocities of sources. Gravitationally bound systems such as galaxies, galaxy groups, or galaxy clusters are expected to follow the Hubble expansion or ‘Hubble flow’. On smaller scales, some inhomogeneity is expected, which results in these systems having some additional, often referred to as ‘peculiar’, velocities with respect to the Hubble flow. Interpreting the CMB dipole as purely a kinematic effect, we can infer for example that our local group of galaxies has a velocity relative to the background of $v = 620 \pm 15$ km/s (Aghanim et al., 2020).

Accounting for the gravitational effects of nearby clusters and superclusters, much of this motion appears to be caused by a yet undetected large-scale inhomogeneity, dubbed the ‘Great Attractor’ (e.g. Dressler, 1991). As more systems are included, the combined peculiar velocity of these systems, or ‘bulk flow’, is expected to become ever smaller at larger scales. For homogeneity and isotropy to hold in Λ CDM, the amplitude of bulk flows has to be limited to certain velocities up to certain averaging volumes. For large enough volumes, the bulk flow velocity of galaxies should be zero, which means they are at rest with respect to the frame of the CMB. Although past results have been consistent with the Λ CDM expectation (Qin et al., 2021, and references therein), recent measurements of bulk flows have suggested that there might be no convergence to the CMB frame, and a rather large bulk flow persists out to distances much larger than what is expected (e.g. Colin et al., 2011; Watkins et al., 2023). Indeed, a convergence not to the CMB frame suggests a difference between the rest frame of the large scale structure and the CMB rest frame, something which would explain the mismatch between the CMB dipole and radio number count dipole. If we are to accept this interpretation of the facts, this violates the cosmological principle upon which our best current cosmological models are built, which can have some deep repercussions.

Working with the assumption that the dipole is caused by a cosmic flow, the standard FLRW metric can include a ‘tilt’ which introduces a preferential direction in the Universe (Ellis and King, 1973). In such a model, density is increased in the direction of the tilt, creating both an intrinsic dipole in the matter distribution, as well as a kinetic dipole as matter naturally moves towards the overdensity. Another possibility is a superhorizon density fluctuation, which could very much induce the same effect as a globally tilted Universe (Domènech et al., 2022). A feature of these models is that different constituents can have different flow velocities, which can even evolve over time. This could be significant, as the CMB dipole represents the radiation component, while the radio number count dipole represents the matter component of the Universe. This degree of freedom would however free up the flow direction for these constituents, making an alignment between these flows extremely unlikely. As these models are further explored, testable predictions should be generated. The superhorizon fluctuation model from Domènech et al. (2022) for instance predicts an additional quadrupole signal in the number counts, which could be pursued as a research venue.

One potentially very interesting effect is that unaccounted for flows can impact cosmological observables, considering that for many past results the CMB rest frame has been considered the true rest frame to which all observations must be transformed before analysis of the data. Such is the case for example for observations of Type Ia supernovae, which are used as standard candles to measure the distance to distant galaxies. Famously, observations of these supernovae showed that the expansion of the Universe is accelerating (Riess et al., 1998; Perlmutter et al., 1999), leading to the inclusion of a cosmological constant Λ in the standard model of cosmology. Tsagas and Kadiltzoglou (2015) show that in a tilted cosmology the measured acceleration can differ from that of the actual Universe. Indeed, Colin et al. (2019) see a dipole in the deceleration parameter q_0 determined from Type Ia supernovae, finding that in this case, the true deceleration can even be consistent with no acceleration at all, which would completely remove the need for dark energy in the standard model of cosmology. Similarly, Mohayaee et al. (2021), when taking large bulk flows into account, show that Type Ia supernova observations favour much smaller values of dark energy density than in standard Λ CDM. Migkas and Reiprich (2018) and Migkas et al. (2021) find anisotropies in several cluster scaling relations that strongly resemble a dipole, although it is misaligned with the CMB dipole. This anisotropy carries through to the measured value of H_0 , which varies across the sky between 66.5 km/s/Mpc and 72.8 km/s/Mpc, which is as big a discrepancy in a single measurement as there currently exists



Figure 1.2: Several individual antennas of the MeerKAT array. Credit: SARAO.

between early time and late time measurements. They suggest that either H_0 truly varies on these spatial scales, or a bulk flow of 900 km/s is present extending out beyond 500 Mpc scales.

The idea that some of the peculiar measurements described above and even some of the other cosmological tensions are all a symptom of the same phenomenon is an enticing prospect. At this point however, the evidence is still uncertain, which makes this an exciting and relatively unexplored frontier of cosmology.

1.4 The MeerKAT Absorption Line Survey

MeerKAT² (Jonas, 2016) is an array of radio antennas located in the Karoo desert in South Africa, consisting of 64 dishes. The antennas are offset Gregorian, with each dish having a diameter of 13.5 meters. Though it can perform different types of observations depending on the science case, for the purposes of this thesis we will focus on interferometric observations, in which all antennas of the array are observing a common target, synthesising a telescope with a diameter equal to the longest baseline for massively improved angular resolution. It saw first light in 2018, and for the first years of operation it has been carrying out predefined large survey projects (LSPs), each focused towards a specific science goal. These LSPs are related to the LSPs and science goals of the Square Kilometre Array (SKA, Dewdney et al., 2009), which MeerKAT will be integrated into once complete.

One of these LSPs is the MeerKAT Absorption Line Survey (MALS, Gupta et al., 2016), which was granted 1655 hours of observing time to search for neutral hydrogen (HI, $\lambda = 21$ cm) and hydroxyl (OH, $\lambda = 18$ cm) absorption lines to map the evolution of cold atomic and molecular gas in galaxies at redshifts $0 < z < 2$. To achieve this, a blind search for absorption lines is conducted by pointing

² Originally the expanded version of the original Karoo Array Telescope (KAT)³

³ Both in Dutch and Afrikaans, 'meer' simply means more.

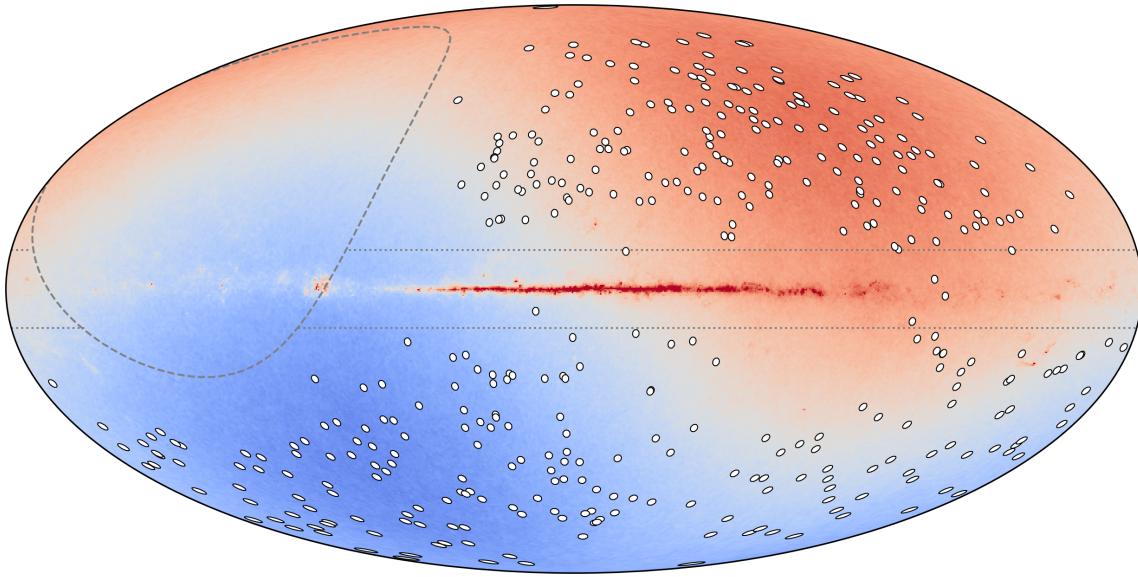


Figure 1.3: Temperature map of the cosmic microwave background at 100 GHz in galactic coordinates, overlaid with 391 MALS pointings observed in L band. The sky area that can not be observed by MeerKAT is indicated by the dashed line, and the area within 10 degrees of the Galactic plane is indicated by the dotted line.

MeerKAT at bright radio sources (Gupta et al., 2022). Due to the high sensitivity and large field of view of MeerKAT, each observation is expected to contain thousands of sources, opening up a plethora of science cases utilising these sources and their properties. With hundreds of pointings observed, the number of sources will be of the order of hundreds of thousands, enough for a dipole measurement. As is shown in Figure 1.3, the observations are spread out across the area of sky observable by MeerKAT, so good coverage of the dipole axis is expected.

Among the surveys with which a dipole measurement has been carried out, MALS will be unique in several aspects. First, rather than a contiguous survey with full coverage of a large part of the sky, MALS provides deep observations of relatively small sky areas. We may expect to detect around 2000 sources per pointing, which would give us nearly 800,000 sources for the 391 pointings of MALS. Though the required number counts for a dipole measurement will be reached, more faint sources will be detected. How this influences the dipole is not known, but any effect or lack thereof may shed light on the origin of the dipole. Second, a common factor hampering the sensitivity of dipole measurements is a lack of intimate knowledge of the systematic effects of the survey. Since dipole measurements are often performed on catalogues without further knowledge of the inner workings of the survey, biases introduced by systematic effects can only be alleviated by making stringent flux density cuts. In the case of MALS, we will be involved in testing and verifying the entire processing pipeline, including calibration, imaging, and source extraction. Each of these processing steps can introduce systematic effects into the catalogue, which we aim to account for, homogenising the catalogues and allowing us to forego stringent flux density cuts. This will ultimately increase the amount of sources that we can use for a dipole measurement, allowing a more sensitive measurement of the cosmic radio dipole.

1.5 This thesis

As described previously, a main hampering issue for measurements of the cosmic radio dipole is systematic effects impacting the homogeneity of the catalogue. This in turn requires strict cuts in the catalogue, reducing sensitivity to a measurement of the dipole. The focus of this thesis is thus to try to alleviate these issues, by accounting for systematic effects rather than trying to remove them from the catalogues entirely. For MALS, this means that we are involved in, follow, and check every step of the way in creating a radio catalogue ourselves, from observation, to calibration, to imaging, and to the cataloguing itself. With this, we envisage to eliminate and account for any systematic effects present in the data, and create a homogeneous catalogue for measuring the cosmic radio dipole. Furthermore, we aim to create improved estimators for the radio dipole that are not only unbiased by the sparse structure of the MALS data, but can also be adapted to account for systematic effects in the data.

This thesis is organised as follows. In Chapter 2, we elaborate the process of producing radio data, from radio measurements to catalogues. In Chapter 3 we present the process of creating homogeneous catalogues of ten MALS pointings. In Chapter 4, we present a number of Bayesian dipole estimators, and test them on simulated and existing radio catalogues. In Chapter 5, we use these concepts on the full MALS dataset of 391 pointings, and perform a dipole measurement. In Chapter 6 we look forward to further steps and improvements that can be made in future to dipole measurements, and we conclude in Chapter 7.

The radio sky: from emission to source counts

Parts of this chapter are based on

Interferometry for dummies

J. D. Wagenfeld

Fundamental Physics in Radio Astronomy research group tutorials, March 2023

Many concepts from this tutorial that have been integrated are inspired by the ‘Essential Radio Astronomy’ textbook by Condon and Ransom (2016) and the ‘Synthesis Imaging in Radio Astronomy’ lecture collection by Perley et al. (1989), however whenever possible I will refer to the original work that those concepts are based on.

Overview

Radio astronomy, in many ways, belongs to a more abstract region of observational astrophysics in which what you see is, in fact, not what you get. If you were to look up at the sky on a reasonably dark night, you would see light from mainly stars, some nearby galaxies, and the Galactic plane. At radio wavelengths, things look vastly different, with the majority of sources emitting strong radio signals being galaxies at extremely large distances. The only nearby star bright enough to see in radio is our own Sun, and while we still see the plane of the Milky way in radio, the origin of this emission is vastly different than at most other wavelengths. A major advantage to observing in radio is that Earth’s atmosphere is transparent to radio waves, and though the sun can be seen at radio wavelengths, it is faint enough that observing during the day is possible as well. Even clouds are transparent to radio waves at GHz frequencies, so observing is possible with most weather conditions.

Creating an image of radio emission is however anything but straightforward. As radio waves have low energy, the methods of capturing them used at other wavelengths do not work. Furthermore, as they have extremely long wavelengths, ranging from centimetres up to metres, spatial resolutions are very poor compared to other wavelengths, even for large telescopes. To reach a reasonable spatial resolution, many telescopes in an array can be combined to ‘synthesise’ a larger radio telescope, which can be done with a process called interferometry. This is a rather complicated process with many caveats and processing steps, which begs close attention if we aim to accurately measure the number counts of radio sources.

2.1 On the nature and distribution of extragalactic radio sources

With ever greater samples of radio sources being gathered, the findings from Ryle and Scheuer (1955) have been expanded into a comprehensive picture of the population of extragalactic radio sources. Though it seems straightforward to simply measure number counts of radio sources, it is important to understand the variety of radio sources that we see in our radio observations. Furthermore, their morphology, spectral shapes, and redshifts all influence a measurement of the dipole. Though there are a great many emission mechanisms that can be observed at radio wavelengths, we are interested in persistent continuum emission. At mid-range radio frequencies ($\nu \sim 1$ GHz), given a limited bandwidth, the spectrum of emission is usually well modelled by a power law. As such, the spectral shapes of sources can with reasonable accuracy be described by

$$S(\nu) \propto \nu^{-\alpha}, \quad (2.1)$$

where α is the spectral index⁴, S the flux density of the emission, and ν the frequency. Though several emission mechanisms exist at these frequencies, the main emission is synchrotron radiation, which is produced by highly relativistic particles (mostly electrons) in magnetic fields. As a consequence, observing the Universe at radio wavelengths reveals the most extreme environments where particles are accelerated to tremendous speeds. The spectral index of sources emitting synchrotron radiation depends on the energy distribution of the relativistic particles, but is usually found to be $\alpha \approx 0.75$ near a frequency of 1 GHz (Condon, 1992).

By far the brightest sources at radio wavelengths, Active Galactic Nuclei (AGN) and the effects they have on their local environment dominate source counts at the brightest flux densities. Due to strong cosmological evolution, the redshift distribution of AGN is sharply peaked at $z \sim 0.8$ (Condon and Ransom, 2016). This is sufficiently far away that we expect these sources to trace the large scale structures of the Universe, and thus abide by the Cosmological principle. As a consequence, the distribution of bright radio sources is expected to be isotropic. Though historically classified into different source types due to their differing morphologies, the currently accepted model of AGN unification (Urry and Padovani, 1995) posits that different observed source types are all products of the same mechanism. The difference in observed phenomena in this case will stem from things such as AGN power, radio loudness, and dust obscuration, and will also largely depend on observational effects such as the viewing angle and observing frequency. In radio continuum observations, two main source types have been historically defined, quasars and radio galaxies. Radio quasars are AGN seen from nearly face on and thus we observe directly the compact AGN core. Much like stars these sources appear point-like in the sky, from which the name quasar, from ‘quasi-stellar radio source’, was initially derived (Chiu, 1964). Seen from wider angles, radio galaxies often appear as extended sources in the sky, where both the core and radio jets are observed. Radio galaxies are usually divided into two subclasses, defined by Fanaroff and Riley (1974). Fanaroff-Riley type I (FRI, Figure 2.1) galaxies usually have lower power, and because most of the emission in the jet originates closer to the core, they are said to be ‘core-dominated’ or ‘edge-darkened’. Fanaroff-Riley type II (FR II, Figure 2.2) galaxies usually have higher power, and because most of the emission originates from the lobes of

⁴The power-law index α is defined like this because for most radio emission its value is usually negative, meaning that flux density increases towards lower frequencies. There is however no consensus on its use, meaning that the inverse definition is also common. For most of this thesis, we will use this definition, as this also means α contributes positively to the amplitude of the number count dipole. There are some exceptions, which I will explicitly point out.




Figure 2.1: Image taken by the Very Large Array of the FRI galaxy 3C 31. Though the radio lobes are visible, most of the emission comes from the region closer to the core. Credit: NRAO.




Figure 2.2: Image taken by the Very Large Array of the FRII galaxy Cygnus A (3C 405). As a typical FRII galaxy, it features a compact core with prominent radio lobes, each lobe sporting a hotspot with increased emission. Credit: NRAO.




Figure 2.3: Image taken by the MeerKAT of the spiral galaxy M83, a known starburst galaxy. Star formation is taking place over the entire galaxy, which is traced by the radio emission. Compact sources of emission within the galaxy can be seen, representing supernova remnants and HII regions, which also indicate recent or ongoing star formation.

the jet, they are said to be ‘lobe-dominated’ or ‘edge-brightened’. In these sources, the emission in the jets originates entirely from the synchrotron process, usually having $\alpha \approx 0.75$, while emission associated with the core can also have a flatter spectrum ($\alpha \lesssim 0.5$). Considering the full population of radio AGN, the spectral index distribution peaks at $\alpha \approx 0.75$, with a tail towards flatter spectra (Condon, 1984; Wilman et al., 2008).

Where AGN dominate the source population at $S \gtrsim 1$ mJy, at lower flux densities we start seeing radio emission from ‘normal’ galaxies that are lacking an active nucleus (Condon, 1992). The radio emission from these galaxies comes from relativistic electrons producing synchrotron radiation in supernova remnants ($\alpha \approx 0.75$), or free-free emission from HII regions ($\alpha \sim 0.1$). The radio emission

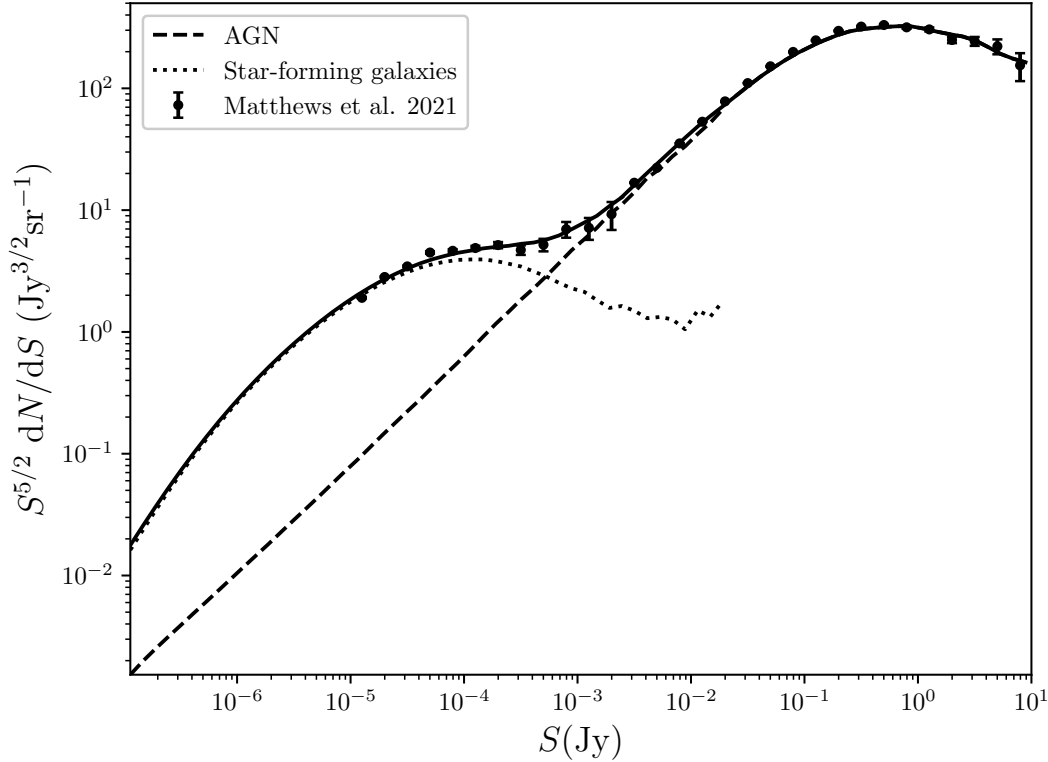


Figure 2.4: Differential source counts for AGN (dashed line), star-forming galaxies (dotted line) and the combined source population (solid line) at $\nu = 1.4$ GHz. The number counts from Matthews et al. (2021b) which currently cover the largest range in flux densities, are also shown (black errorbars).

of these galaxies is thus an excellent tracer of recent star-formation, and if the galaxy is resolvable this might even be used to track where in these galaxies the star-formation is taking place. Therefore, it is common to refer to these galaxies as ‘star-forming galaxies’ (SFGs, Figure 2.3). However, at the cosmological distances that we are interested in for a dipole measurement, most SFGs are exceedingly faint and are usually unresolved. At present, radio surveys with large sky coverage have not been deep enough to probe the population of SFGs, so not much is known about the distribution of these sources.

The way that these source populations are usually studied in the context of cosmological observations stems from the earliest observations in which radio sources were used to constrain cosmological evolution. This involves looking at the number counts of radio sources at different flux densities, where we measure the differential number of sources dN per steradian between flux densities S and $S + dS$, where dS is the size of the flux density bin. As mentioned in Chapter 1, the expectation of a static Universe is that the number of radio sources follows $S^{-3/2}$, such that the number of sources per unit of flux density dN/dS would be proportional to $S^{-5/2}$. As such, the differential number counts of radio sources are usually multiplied by $S^{5/2}$, which would yield flat relation with respect to flux density in the case of a static Universe. Of course, it has been long proven that this is not the case, and taking into account the source populations and their cosmological evolution yields a curve such as the one

shown in Figure 2.4. Although observations have not been deep enough to probe the full population of sources, we can predict the expected number counts of the populations with simulations, such as the SKA Design Study (SKADS) simulated skies (Wilman et al., 2008) or the Tiered Radio Extragalactic Continuum Simulation (T-RECS, Bonaldi et al., 2019). To get the most complete picture of number counts at different flux densities, shallow observations covering a large area can be combined with deep observations covering a small area (e.g. Matthews et al., 2021b). Now that the deepest surveys are probing into the population of SFGs, the differential number counts of faint sources can be used to constrain the star-formation history of the Universe as well (e.g. Mauch et al., 2020; Matthews et al., 2021a). Finally, differential number counts are usually the most convenient method of comparing different catalogues of radio sources and comparing observations to simulations when we have large catalogues containing many sources.

2.2 Radio observations and synthesis imaging

In many aspects, that we are able to produce images from radio data can be considered a feat in itself. Considering that this process consists of many steps, includes a number of uncertainties and assumptions, it is not difficult to imagine that systematic effects are easily introduced in the process. In terms of dipole measurements, this argument is often used to scrutinise the anomalous results, making a solid understanding of the process extremely important.

At the frequencies we are currently interested in ($\nu \sim 1$ GHz), radio waves are captured by reflecting them off of a parabolic dish into a receiver. Changing the size of the dish will not only change the amount of light captured, but will also change the angular resolution. The sensitivity pattern of a single dish can be described by the Fourier transform of the illumination pattern of the dish. The top plots of Figure 2.5 show the most simple case, a uniformly lit dish, where the sensitivity pattern, commonly referred to as ‘primary beam’, follows a squared sinc function that peaks at the location the dish is pointing at, the pointing centre. The main lobe of the beam corresponds to the central peak of this function, while sidelobes correspond to peaks of the function away from the centre. These sidelobes are undesirable as they can pick up on emission of sources that the telescope is not directly pointing at. They can be suppressed by changing the aperture illumination. The MeerKAT dishes use a cosine tapered illumination function (e.g. Condon and Ransom, 2016; Mauch et al., 2020), shown on the bottom plots of Figure 2.5, which suppresses sidelobes. In either case, the full width at half maximum (FWHM) of the main lobe determines the angular resolution θ , which is related to the size of the dish as

$$\theta = 1.02 \frac{\lambda}{D}, \quad (2.2)$$

where λ is the wavelength of the light, D the diameter of the dish.

Because radio waves have relatively long wavelengths ($\nu = 1$ GHz corresponds to $\lambda = 30$ cm), radio dishes have to be much larger than telescopes operating at other wavelengths to get a good angular resolution. In addition to this, a single dish usually focuses all light towards one receiver, meaning one pointing of the telescope essentially gives a single pixel on the sky. Images can thus only be obtained by scanning across the sky with the telescope, which is a time-consuming process, and even for extremely large dishes the resolution will be extremely coarse compared to telescopes operating at shorter wavelengths.

Better resolution can be achieved by combining multiple radio dishes, pointing them at the same

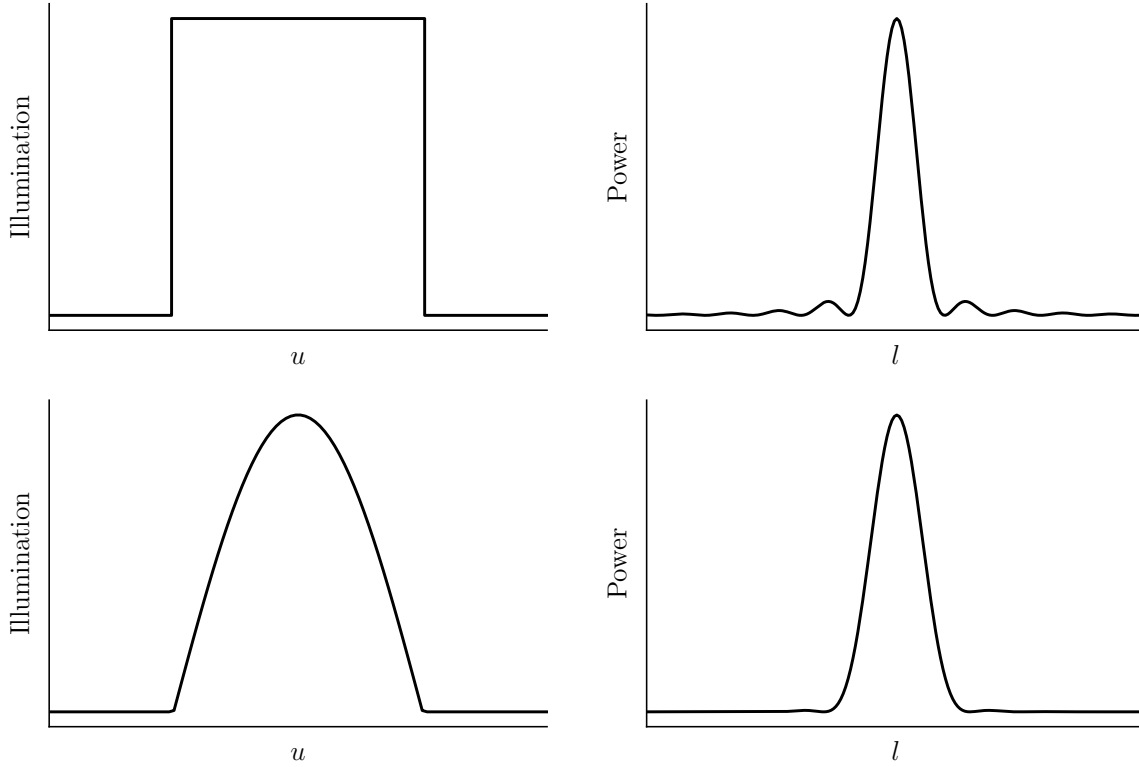


Figure 2.5: Aperture illumination functions (left) and the corresponding primary beam patterns (right), which are the square of the Fourier transforms of the illumination function. Compared to uniform illumination (top), the cosine tapered illumination (bottom) better suppresses sidelobes of the primary beam.

location in the sky, and accounting for the geometrical distance between the dishes. Figure 2.6 shows a schematic representation of the simplest of such interferometers, having only two elements, and the resulting signal. Cross-correlation of the signals received by the dishes eliminates all signals that are not common between them, which also removes noise patterns of the individual antennas. The sky signal common between them gets picked up, and as the interferometer will pick up any signal that is in phase between the two dishes, the sensitivity pattern follows a sinusoidal form on the sky. In interferometry this pattern is commonly referred to as a ‘fringe’. The width of the fringe is determined by the wavelength λ and the baseline length B , which describes distance between the dishes,

$$\theta = \frac{\lambda}{B}. \quad (2.3)$$

As shown in Figure 2.6, if observing a source with some arbitrary brightness distribution $I(\hat{s})$ on the sky, the signal received by the interferometer is equal to the emission, multiplied by the fringe pattern, and integrated over the covered sky area $d\Omega$,

$$R_c = \int I(\hat{s}) \cos(2\pi\nu\vec{B} \cdot \hat{s}/c) d\Omega. \quad (2.4)$$




Figure 2.6: Schematic representation of a two element interferometer with baseline \vec{B} observing a common source in the direction \hat{s} . The incoming waves have a common field strength E and are oscillating with frequency $2\pi\nu = 2\pi c/\lambda$. The difference in path length that the light has to travel introduces a phase difference between the signals picked up by the antennas, given by the geometric time delay, $\tau_g = \vec{B} \cdot \hat{s}/c$. The incoming signals are multiplied (\times) and then time averaged ($\langle \rangle$) by the correlator, which yields the fringe pattern of the baseline, with power $P = E^2/2$ and phase $2\pi\nu\tau_g$. The aforementioned correlation produces a cosine fringe pattern R_c , and a sine fringe pattern R_s can be produced by shifting the phase of one of the signals by 90° before correlation. A correlator producing both R_c and R_s signals is referred to as a complex correlator.

The cosine fringe R_c will only pick up the emission coinciding with the positive peaks of the fringe, or the ‘even’ part of the emission. To pick up on the ‘odd’ part of the emission, a sine fringe R_s needs to be produced as well. As shown in Figure 2.6, a complex correlator produces both a sine and cosine fringe, which combines the signals as a complex exponential using Euler’s formula (Euler, 1748),

$$e^{i\phi} = \cos(\phi) + i \sin(\phi), \quad (2.5)$$

This defines the signal seen by the interferometer, the complex visibility \mathcal{V} , as

$$\mathcal{V} = R_c - iR_s = A e^{-i\phi}, \quad (2.6)$$

where A is the visibility amplitude and ϕ is the visibility phase. Following Equation 2.4, the relation




Figure 2.7: (u, v) -coverage of a MeerKAT observation of an hour of observing times, carried out in three twenty minute blocks.

between the complex visibility and an arbitrary brightness distribution on the sky is then

$$\mathcal{V} = \int I(\hat{s}) \exp(-i2\pi\vec{B} \cdot \hat{s}/\lambda) d\Omega, \quad (2.7)$$

which is a Fourier transform.

While a single pair of dishes does a poor job of locating the signal on the sky, adding more baselines with different lengths will modify the sensitivity pattern to suppress sidelobes while leaving the central peak intact. Though the simplest interferometer consists of two dishes, any interferometer with N dishes can be treated as $N(N-1)/2$ two-dish interferometers. With more and more baselines, the sensitivity pattern, also referred to as the synthesised beam or point-spread-function (PSF), will approach a Gaussian with a FWHM given by Equation 2.3 for the largest baseline in the interferometer. With such an interferometer, we essentially simulate having a much bigger dish, with a diameter equal to the longest baseline. Of course with such a setup, there will be many gaps in the data, which will be reflected in the shape of the synthesised beam. While the resolution of the interferometer is determined by the longest baseline, while the maximum observable scale of emission is determined by the shortest baseline.

To describe the situation quantitatively, we define a coordinate system in which the baselines live, which are (u, v, w) coordinates, which describe the baselines of the array essentially as seen from the observing target. The corresponding coordinate system in the sky is (l, m, n) . Extending Equation 2.7 to this description, the intensity on the sky is related the visibility via

$$\mathcal{V}(u, v, w) = \int \int \frac{I(l, m)}{n} \exp[-2i\pi(ul + vm + wn)] dl dm. \quad (2.8)$$

The observed sky area is here $d\Omega = dl dm/n$. This equation is however not a Fourier transform, which will create problems with imaging. A Fourier transform can be obtained from this equation by making the assumption that w , the axis in the direction of the target on the sky, is equal to zero, which means that the entire array must lie in a single plane. With this, Equation 2.8 can be reduced to

$$\mathcal{V}(u, v) = \int \int I(l, m) \exp[-2i\pi(ul + vm)] dl dm. \quad (2.9)$$

This is a two-dimensional Fourier transform, describing the measured values at each point on the (u, v) plane. Each baseline occupies the two points on the (u, v) plane which describe the vector of the baseline. As we observe with the array for an extended period of time, the earth rotates, thereby incrementally changing the positions of the baselines as seen from the observing target. This fills up the (u, v) plane, improving the (u, v) -coverage in a way that is much more efficient than adding additional dishes to the array. This technique is called rotation aperture synthesis, and is the most commonly used method of improving the (u, v) -coverage of observations. Figure 2.7 shows the (u, v) -coverage after an hour of observing with MeerKAT, showing the tracks on the (u, v) plane made by the baselines over time.

Observing with an interferometer yields the visibility values at all measured points in (u, v) space, and as such these data points, which comprise the raw data of radio observations, are usually referred to as visibilities. The (u, v) -coverage determines what scales the interferometer is sensitive to, relating directly to the synthesised beam through a two-dimensional Fourier transform. Figure 2.8 shows the relation between the different fields, showing that the raw or ‘dirty’ image is the inverse Fourier transform of the measured visibilities and as that is a multiplication of the sky visibility with the (u, v) -coverage, the dirty image is the convolution of the true sky intensity with the synthesised beam.

The Fourier transforms required to get an image from the visibilities are not performed analytically, but rather numerically through fast Fourier transforms (FFTs). This requires discretisation of the data, in particular the visibility data in the (u, v) plane. When different (u, v) samples fall onto the same pixel, the question is how to weight them when combining them into a final value of the pixel. Natural weighting gives equal weight to all samples, thus weighing the pixels where more samples are present more heavily. Since usually the core of the array is usually more heavily sampled, as is also seen in Figure 2.7, the corresponding larger scales are emphasised. This can create unwanted structures in the synthesised beam, but will yield the best sensitivity. In contrast, uniform weighting weighs all pixels according to the amount of samples, thus weighing each pixel equally, which will emphasise the smaller structures seen by long baselines. This weighting yields the best resolution, but worse sensitivity. The best solution usually lies somewhere in between, which is offered by Briggs weighting (Briggs, 1995). As shown in Figure 2.9, Briggs weighting has a tunable robust parameter R that can produce weighting anywhere between natural ($R = 2$) and uniform ($R = -2$) weighting.

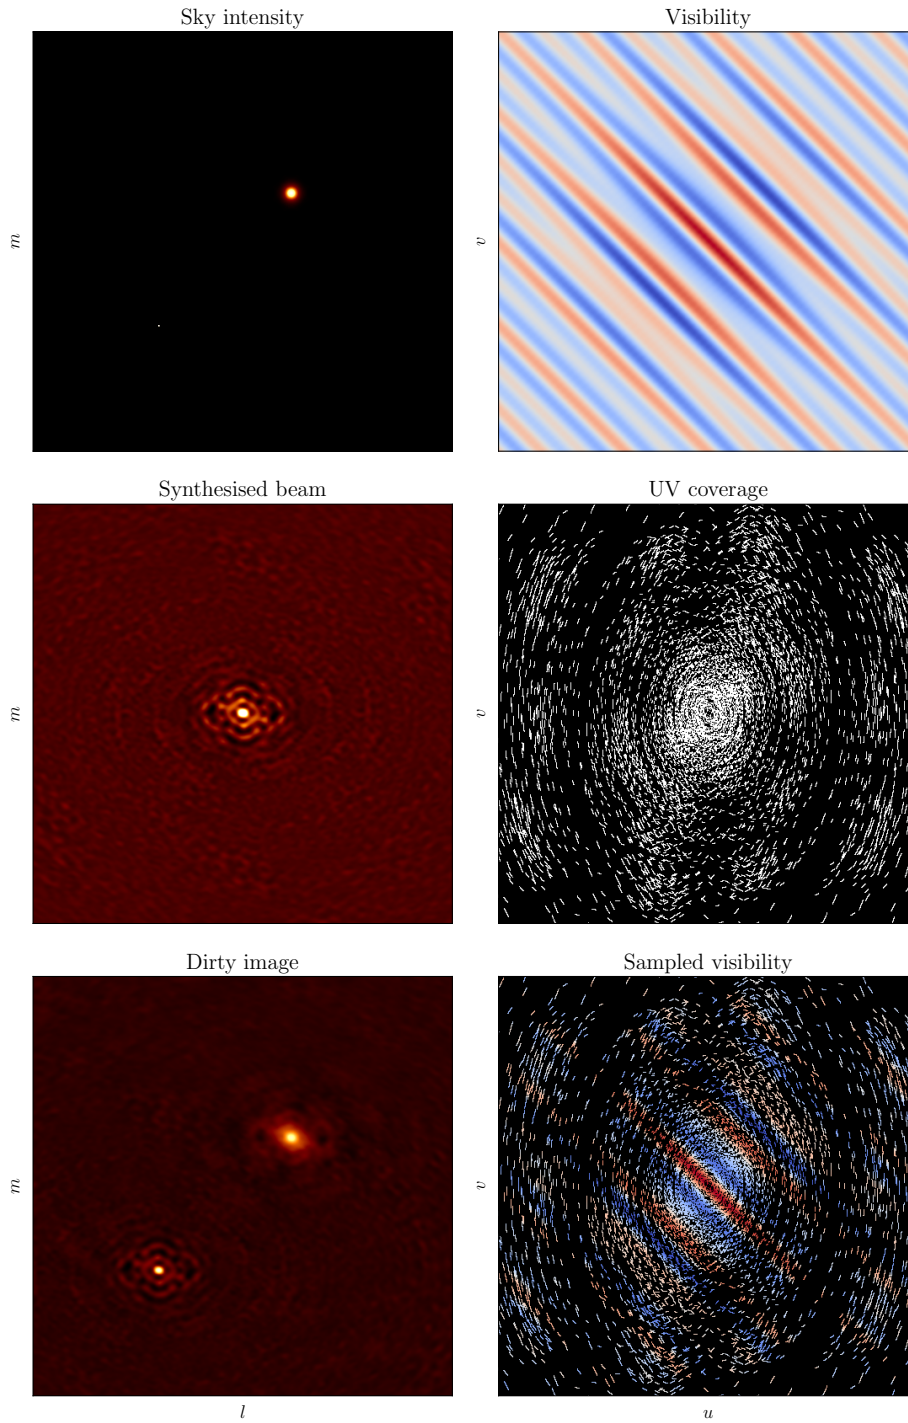


Figure 2.8: Definitions of images in interferometry. In each case the left image can be transformed to the right image through a two-dimensional Fourier transform, and the right image can be transformed to the left image via an inverse two-dimensional Fourier transform. From left to right, top to bottom: (i) Sky intensity, (ii) Sky visibility, (iii) Synthesised beam or PSF, (iv) UV-coverage, (v) Dirty image, (vi) Sampled sky visibility.

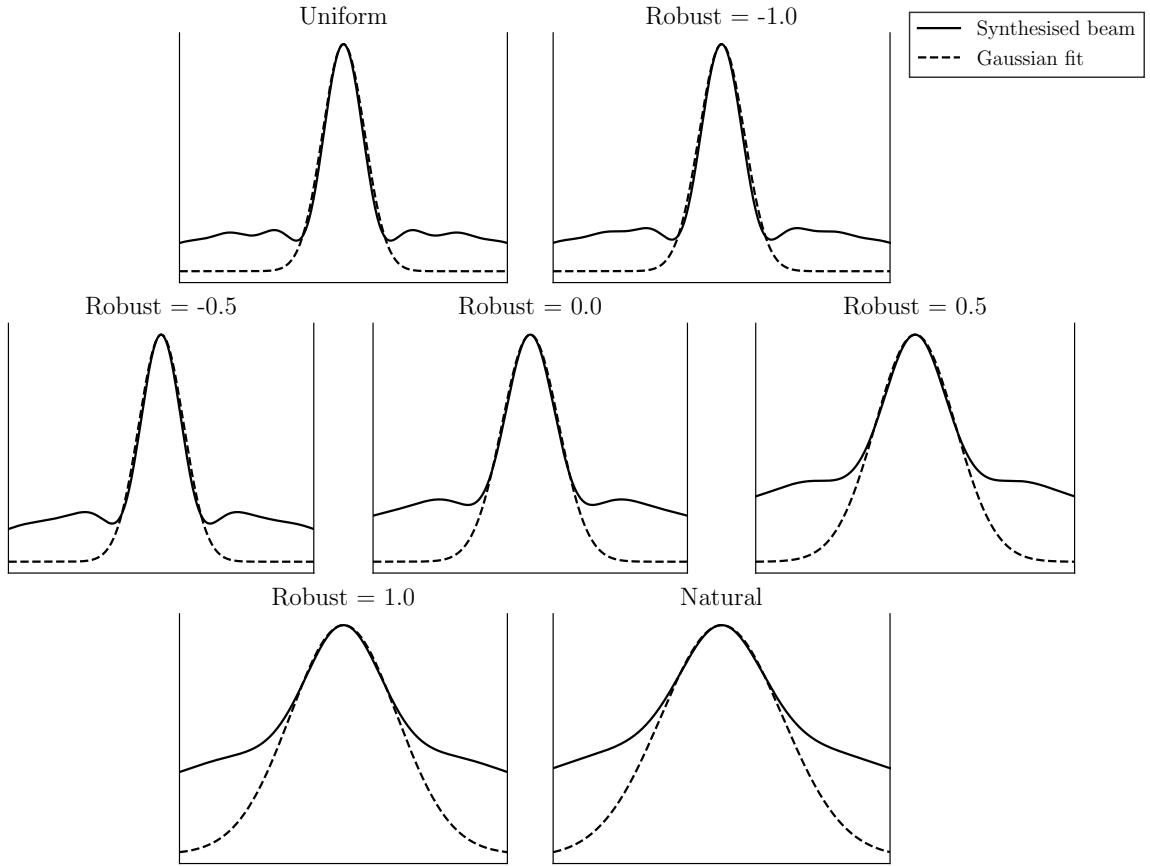


Figure 2.9: Cross-cut of the synthesised beam (solid line) and a Gaussian fit to the beam (dashed line) for different weightings of MeerKAT UV-coverage, going from uniform ($R = -2$) to natural ($R = 2$). For high robust values and natural weighting the core of MeerKAT is highly emphasised, adding wings to the synthesised that are not well fit by a single Gaussian.

2.2.1 Observation and calibration

In practice, radio receivers have two feeds that are meant to pick up not just the amplitude of the radio wave, but also the polarisation, which describes the orientation of the incoming wave. The most common feeds in radio telescopes are either circular (R, L) or linear (X, Y) feeds. MeerKAT uses linear feeds, however the treatment is not much different for circular feeds, and for science purposes the signals from these feeds are always converted to the universal Stokes visibilities I, Q, U , and V (Stokes, 1851). For a single baseline between antennas i and j , there are thus four signals, two from each of the two antenna feeds, and a visibility can then be represented as a vector contained the cross-correlated products of the feeds,

$$\vec{\mathcal{V}}_{ij} = \begin{pmatrix} V_{X_i} V_{X_j} \\ V_{X_i} V_{Y_j} \\ V_{Y_i} V_{X_j} \\ V_{Y_i} V_{Y_j} \end{pmatrix}. \quad (2.10)$$

The Stokes visibilities can also be represented as a vector, and can be obtained from the above cross-correlation products by

$$\begin{pmatrix} I \\ Q \\ U \\ V \end{pmatrix} = \begin{pmatrix} V_{X_i}V_{X_j} + V_{Y_i}V_{Y_j} \\ V_{X_i}V_{X_j} - V_{Y_i}V_{Y_j} \\ V_{X_i}V_{Y_j} + V_{Y_i}V_{X_j} \\ i(V_{X_i}V_{Y_j} - V_{Y_i}V_{X_j}) \end{pmatrix}. \quad (2.11)$$

For the purposes of source counts and a dipole measurement, we will only focus on Stokes I , which represents the emission from all polarisations combined, and thus measures the full intensity⁵.

Before creating a radio image, there are a number of physical and instrumental effects which impact the radio emission and must be accounted for. A common way to describe these effects is through cumulative corruptions on the ‘perfect’ visibilities to produce the observed visibilities in the Hamaker-Bregman-Sault Measurement Equation (Hamaker et al., 1996; Sault et al., 1996);

$$\vec{\mathcal{V}}_{ij}^{obs} = \vec{J}_{ij} \vec{\mathcal{V}}_{ij}. \quad (2.12)$$

Here there is one visibility $\vec{\mathcal{V}}_{ij}$ for each combination of baseline between antenna i and j , sampling time, and spectral channel. \vec{J}_{ij} represents the accumulation of all effects on these visibilities and can be refactored to isolate all the individual effects. As mentioned, many such effects exist, but one way to refactor the measurement equation is as follows:

$$\vec{J}_{ij} = \vec{K}_{ij} \vec{B}_{ij} \vec{G}_{ij} \vec{D}_{ij} \vec{E}_{ij} \vec{P}_{ij} \vec{F}_{ij}. \quad (2.13)$$

Each component of the equation is a Jones matrix representing one effect on the incoming radiation, and the order of the matrices actually represents the order in which the physical effects occur. In this description, the terms have the following meanings:

- F Faraday rotation
- P Parallaxic angle rotation of the antenna feed with respect to the sky
- E Primary beam effects
- D Polarisation leakage between the antenna feeds
- G Electronic gain response
- B Bandpass or frequency response
- K Geometrical delay and delay-rate between antennas

Though many such effects exist, we focus on the effects that have an impact on continuum data, and only consider those that impact Stokes I and forego the potential effects on the measurement of polarised emission. In order to correct the data for these effects, several calibrators are usually observed along with the target. A flux calibrator is usually an extremely bright, non-variable source, of which the flux density and spectrum is well known. It is used to define the flux density scale of the observations, as well as determine the bandpass or frequency dependent (B) response of the telescope.

⁵ The prospect of measuring the dipole in polarisation is touched upon in Section 6.2




Figure 2.10: Observational setup of an average MALS observing run. A bright flux calibrator is observed at the start, the end, and once in between for ten minutes. The three targets are observed in alternating fashion in twenty minute blocks, for a total integration time of an hour per target. Before and after each target observation a nearby gain calibrator is observed for one minute.

It is usually observed once at the start and end of an observing run. Additionally, a gain calibrator is a reasonably bright source that is close on the sky to the target. As the name suggests, a gain calibrator is used to correct the gain response (G) of the electronics in the system. It should be bright enough that a short observation is enough to reliably determine these solutions. As the gains also vary in time, it is usually observed shortly before and after observations of the target source. The most common observing patterns for total intensity science include a flux calibrator and gain calibrator.

A typical observation of MALS is shown in Figure 2.10, following the previously described observing pattern. With this information, the data can be calibrated for the bandpass (B) and time variable gain (G) response of the antennas. For G the amplitude and phase, as defined in Equation 2.6, vary on different timescales, with phase usually varying on shorter timescales. These are thus solved for separately with different solution intervals, i.e. different timescales over which the data is averaged and fitted for. Finally, the solutions obtained from the calibrator sources are applied to the target to create calibrated visibilities that can then be imaged.

2.2.2 Imaging

Figure 2.8 shows that performing a Fourier transform on the visibilities produces a dirty image, which is the sky intensity convolved with the synthesised beam. The resulting image and synthesised beam also depends on the chosen weighting when discretising the (u,v) samples. This is technically the most correct representation of the data, given the incomplete sampling of the (u,v) plane. What we




Figure 2.11: Original dirty image from Figure 2.8 (left) and the same image after Högbom CLEAN (right).

want however is an accurate representation of the sky intensity, where we can be confident that all we see in the image represents real radio emission. To achieve this, we need to create a good model of the emission in the field.

If we were to know nothing about the emission we were receiving there would be no way to differentiate between real emission and image artefacts, however there are certain rules that are followed by nature, which we can use to constrain what the emission is supposed to look like. Creating a model of the sky will then allow us to disentangle the real emission from the effects of the synthesised beam, which means we can deconvolve the image. This process is usually performed during imaging of the field, and is most commonly referred to as ‘CLEAN’⁶. The most basic but ubiquitous CLEAN algorithm was introduced by Högbom (1974), and goes as follows:

1. Find the brightest pixel in the dirty image
2. Subtract from the dirty image, at the position of this pixel, the dirty beam multiplied by the brightness of the pixel, times a gain factor
3. Add the subtracted value at the position of the pixel to the model image
4. Repeat steps 1-3 until the brightest pixel is below some user-specified threshold, or a user-specified number of iterations is reached
5. Convolve the model image with the CLEAN beam, which is usually an elliptical two-dimensional Gaussian fitted to the central lobe of the of the dirty beam, to create the CLEAN image
6. Add the CLEAN image to the residual dirty image to produce the final image

⁶ Even though it is written in capital letters, it does not appear to be an acronym for anything.

If CLEAN is performed well, an image is produced that appears as shown in Figure 2.11. The emission is deconvolved with the synthesised beam, removing the artefacts in the image, and convolved again with the CLEAN beam, essentially appearing as if no gaps in the (u,v) -coverage were present at all.

The Högbom CLEAN algorithm is rather simplistic, in that it assumes all emission on the sky can be modelled by a combination of point sources. This was a reasonable assumption for older radio telescopes which had rather poor resolution and thus could not resolve many sources. It also does not take into account the frequency structure of the emission, or any effects which occur away from the pointing centre. There are however many extensions to CLEAN that have been implemented in order to account for these effects. In the case of MeerKAT, the compact core of the array is especially sensitive to emission on larger scales, so a multi-scale algorithm must be used to model this emission. Furthermore, MeerKAT has a large bandwidth ($\Delta\nu \sim 800$ MHz), which necessitates the use of multi-frequency algorithms which model the emission as a function of frequency as well. An example of such an algorithm is multi-term multi-frequency synthesis (MTMFS, Rau and Cornwell, 2011), which models the emission at multiple scales, and models the spectrum of the emission with a Taylor series expansion

$$\vec{I}_\nu = \sum_{t=0}^{N_t-1} w_\nu^t \vec{I}_t \quad \text{where} \quad w_\nu^t = \left(\frac{\nu - \nu_0}{\nu_0} \right)^t. \quad (2.14)$$

Here N_t represents the order of the Taylor expansion, and ν_0 the central frequency of the band. Besides the regular image, this can also yield a spectral index (α) image, and even a spectral curvature (β) image if modelling with two and three Taylor terms respectively, such that the emission is can be characterised as

$$I_\nu = I_{\nu_0} \left(\frac{\nu}{\nu_0} \right)^{\alpha + \beta \log(\nu/\nu_0)}. \quad (2.15)$$

If we are interested in the full field of view of the telescope, as we are if we want to get as many sources as possible, there are additional effects that are introduced away from the pointing centre. As mentioned shortly before, the two-dimensional Fourier transform which we defined in Equation 2.9 does not hold if w is nonzero for the baselines with which we are observing. With a larger field of view, the w becomes nonzero due to us observing the celestial sphere which has an extra dimension compared to the plane of the sky which we can assume for small field of view. As such, the simple Fourier transform of Equation 2.9 is no longer valid, which will produce errors and artefacts in the image. One way to deal with this is with faceting, in which the full field of view is separated into multiple facets, with each facet having its own phase centre. As the w -term is essentially a phase offset, this alleviates the errors stemming from this offset. Another way is to use w -projection, in which instead of imaging in a single plane the data is imaged in many planes, each at a different value of w . For each single plane Equation 2.9 is valid again. A further effect which affects the imaging with large field of view are variations caused by things such as pointing errors, parallactic angle rotation of the feeds, and different dish sizes. These affect the aperture illumination functions of the individual antennas, which can vary in time, frequency and polarisation. These variations carry on through to the primary beam of the telescope as a function of the aforementioned quantities. They can be accounted for using A-projection.

2.2.3 Self-calibration

The calibration process described before produces calibrated visibilities for the target, which should be reasonably close to the ‘perfect’ visibilities, and can thus be imaged. Usually however, the gains can vary on timescales shorter than the observing time of the target. These timescales are too short to solve for with the gain calibrator, however if there are sources in the target field with high enough signal to noise, we can use the target field itself to perform additional calibration. This is, not surprisingly, often referred to as self-calibration. Though the observation of a gain calibrator is always centred on the calibrator itself, the structure of the target field in any given observation is most likely not a priori known. Thus, in order to perform this calibration, we need to have a good model of the (bright) emission in the field.

It is not difficult to see that the CLEAN procedure synergises quite well with self-calibration, as it naturally produces a model of the (real) emission in the image. To enable self-calibration, the image model derived during CLEAN can be Fourier transformed to produce model visibilities. These in turn can be used to perform gain calibration, which is usually done in an iterative fashion. A CLEAN run is followed by a round of gain calibration, after which another round of CLEAN is performed, which is followed by another round of gain calibration, and so forth. Each iteration of this uses a shorter solution interval for the gain calibration, and CLEANs deeper (i.e. with more iterations or to a lower threshold). Multiple iterations are done for phase-only calibration, each time shortening the solution interval. Signal to noise permitting, self-calibration is concluded with a round of amplitude calibration with a long solution interval, followed by a deep CLEAN run. In the case of MALS, the central source is always bright enough to allow self-calibration, so all steps of self-calibration can be carried out to improve the image.

2.3 Image analysis

After calibration, self-calibration and imaging, we are left with a science image where the morphology inside the image represents the emission of the sky as best as the data and the allocated processing time allows. The bright emission is convolved with the CLEAN beam, represented by a two-dimensional Gaussian. Weaker emission below the CLEAN threshold, including the noise, is still convolved with the dirty beam. The sensitivity of the resulting image is multiplied by the primary beam (as shown in Figure 2.5), which decreases further away from the pointing centre. For an accurate measure of source counts in an image, these effects have to be taken into account.

2.3.1 Primary beam correction

The effect of the primary beam on a radio image is that the true sky intensity is multiplied by the primary beam pattern. This means further away from the pointing centre, sources are measured with less flux density than they have in reality, and due to the noise increase less sources are detected. While we cannot recover the sources that are below the sensitivity limit due to the primary beam response, we must correct for the primary beam response to get correct flux densities for the radio emission in the image. To compensate for the primary beam response, we need an accurate model of the primary beam of the telescope.

One way to obtain a model of the primary beam is to derive it analytically, which can be done given the illumination pattern of the telescope. As described previously, the primary beam response can




Figure 2.12: Cross-sections of primary beam models, taking into account only the central frequency (dashed line) or the full band, as determined by fitting a Taylor polynomial to different primary beams within the full band of MeerKAT (solid black line). This is done for full intensity (left) and spectral index (right), with the beams at individual frequencies determined with holographic measurements. The wideband primary beam for full intensity is slightly broader than its single frequency counterpart, while the spectral index primary beam can not be determined at all with a primary beam at a single frequency. Although minor, some directional asymmetries can be seen.

be described by the square of the Fourier transform of the illumination pattern. For Stokes I, such a primary beam model can be assumed to be directionally symmetric, an idealisation of the actual primary beam. In reality there are several factors that can influence the shape of the primary beam. Deformations in the dishes of the individual antennas, which can be influenced by environmental factors such as the temperature, change their individual primary beam responses. This carries on through to the primary beam of the entire telescope, introducing variations in the variation of the primary beam that are not directionally symmetric. As mentioned previously, these effects can even vary as a function of time, frequency, and polarisation, which must be addressed using A-projection. The most precise way to determine the primary beam response of a telescope is through measurements using holography, but even this will only present an average, and will not account for the small differences between observations.

A further complication is that the primary beam is frequency dependent. This is especially relevant when a large bandwidth is used for a continuum image, as emission is multiplied with different primary beams at different frequencies. The wideband primary beam can be reconstructed by applying the same multi-frequency synthesis method as was applied on the image itself. This requires a set of primary beams with the same frequency structure of the images. The way the primary beams of the

individual frequencies are combined depends on the employed multi-frequency method. Figure 2.12 shows primary beam models which were determined using the MTMFS method as described in Equation 2.14, fitting a Taylor polynomial to primary beam models at individual frequencies over the full band. Using MTMFS, the primary beam response of the full band can be calculated as the zeroth order Taylor terms, as well as the higher order terms which can yield the beam response in terms of spectral index.

2.3.2 Source extraction and cataloguing

Once images became large and deep enough to contain hundreds or even thousands of sources, finding and characterising these sources in a systematic way is only feasible in an automated fashion. The principle of these is simple: measure the mean value μ and standard deviation σ across the image, and find emission peaks where the flux density $S > \mu + f\sigma$, with f some given threshold value. For reliable results, $f = 5$ or higher is usually chosen. A number of general source finding algorithms are available, however for radio astronomical images some more specialised algorithms exist. In radio, the peaks are fit with two-dimensional Gaussians, as most sources are likely to be close in terms of shape to the Gaussian representing the clean beam of the image.

For an accurate measure of number counts and by extension the dipole, there are a couple of measures that are extremely important to assess with regards to source extraction and the produced catalogues. One aspect is the completeness of the catalogue, which is a quantity that describes what fraction of sources in the sky is actually detected in the image. This fraction is usually dependent on local noise, which can vary across the survey, and flux density of the sources. Probably less impactful, but still important to assess, is the purity of the catalogue. Purity quantifies what fraction of detected sources are real, as there is always a small chance that some noise fluctuations and artefacts around bright sources are strong enough to be picked up by source finding. Both these measures require intermediate data products to determine, which are usually not publicly available for surveys. In the next chapter, we will assess these quantities along with the full calibration and imaging pipeline of MALS.

The most homogeneous of catalogues

The MeerKAT Absorption Line Survey: Homogeneous continuum catalogues towards a measurement of the cosmic radio dipole

J. D. Wagenveld, H.-R. Klöckner, N. Gupta, P. P. Deka, P. Jagannathan, S. Sekhar, S. A. Balashev, E. Boettcher, F. Combes, K. L. Emig, M. Hilton, G. I. G. Józsa, P. Kamphuis, D. Y. Klutse, K. Knowles, J.-K. Krogager, A. Mohapatra, E. Momjian, K. Moodley, S. Muller, P. Petitjean, P. Salas, S. Sikhosana, and R. Srianand
Astronomy & Astrophysics, 673, A113 (2023)

Overview

In this chapter, we assess the processing pipeline of MALS, from imaging to cataloguing, the steps of which have been described in the previous chapter. This analysis was performed on ten MALS pointings, focusing homogenising the catalogue in preparation of a dipole measurement.

The calibration, self-calibration, and imaging of the data was carried out by the Automated Radio Telescope Imaging Pipeline (ARTIP, Gupta et al., 2021) used for MALS data. To assess the process, we evaluated the logs that were generated by the pipeline, extracting relevant information and checking error messages. From the imaging products, we additionally produced spectral index images. To correct for the primary beam attenuation in the images, we created and applied wideband primary beam models which take into account the fact that we are imaging the full bandwidth.

We performed source finding on all images with the Python Blob Detection and Source Finder (PYBDSF, Mohan and Rafferty, 2015). The products created by PYBDSF included source catalogues, which were modified to include additional columns, and root-mean-square (rms) maps, which we used to assess the noise properties of the images. To quantify the noise structure of the pointings, we defined the rms coverage as the cumulative distribution of pixel values in the rms map. Due to effects of the primary beam and the fact that a strong source is always present at the centre of each pointing, we found that the noise structures in all of the pointings are very similar. We were able to characterise the differences between the pointings by a single rms value at a rms coverage of 20%, and define this value as σ_{20} .

To assess how well Gaussian components are combined into sources, we created cutouts of the sources with multiple components, visually inspected them, and assigned quality flags. We created

mock observations by injecting sources from the SKADS simulations (Wilman et al., 2008) into the residual images produced by PyBDSF. With these mock observations, we assessed the completeness of the catalogues as a function of flux density, separation from the pointing centre, and source size. We used the same mock observations to check the flux recovery statistics of the catalogues, and we finally measured the purity of the catalogues. We first performed these checks on the individual pointings, and then on the combined catalogue of all the pointings. We were able to combine completeness and flux recovery measures using the previously defined σ_{20} for a universal measure of these quantities.

From these ten pointings, we produced a combined catalogue containing 16,313 sources and covering 37.5 square degrees of sky down to a sensitivity of 10 $\mu\text{Jy}/\text{beam}$. In order to check potential systematics in flux density and astrometric precision, we cross-matched the combined catalogue with the RACS and NVSS catalogues. We investigated different source populations in the catalogue, checking how sources with multiple (disconnected) components might influence source counts. To assess the quality of spectral indices produced during imaging, we compared them with the spectral indices of simulated sources from SKADS, as well as spectral indices measured by comparing the flux densities between different individual spectral windows from the first data release of MALS (MALS DR1, Deka et al., 2023).

Finally, we produced differential source counts, and corrected them using the rms coverage and completeness measures of the pointings. These corrections were performed on a simulated data set as well as the real data. We found that these corrections can be used to account for the incompleteness of the catalogue, producing corrected number counts down to 300 μJy . With this, we expect that we can homogenise the catalogues and properly account for systematic effects once the full set of MALS pointings is available, which will have enough sources for a dipole measurement.

This paper was originally published as Wagenveld et al. (2023a) in *Astronomy & Astrophysics* in 2023. It can be found in its original form in Appendix B. It is reproduced below with some minor adjustments to support the flow of this thesis. As the lead author of the paper, I wrote the manuscript and produced all the figures. I produced the catalogue that was released along with the paper, and performed the checks and analysis on the data.

3.1 Introduction

Though the expected MALS number counts are sufficient for a dipole measurement, a dipole estimate requires a homogeneous catalogue. Systematic effects influencing the sensitivity of surveys are common and are usually dealt with by making conservative cuts in the data to avoid biasing the dipole estimate. Instead, in this work we present a thorough analysis of ten MALS pointings, aiming to fully understand the systematics present in the survey data. This will allow us to account for these systematics when measuring the radio dipole using hundreds of MALS pointings. The nature of the survey provides additional challenges for this type of measurement. Previously, measurements of the dipole have been performed with contiguous surveys such as the NVSS, whereas MALS will be sparser, sampling the sky in many different directions. However, compared to these surveys, MeerKAT has a much higher sensitivity (10 $\mu\text{Jy}/\text{beam}$), which allows us to probe deeper into the population of faint radio sources. Furthermore, past dipole measurements from contiguous sky surveys have been performed post-factum, with little knowledge of the internal processing, and therefore present systematics of these surveys beyond what is described in the literature. In this paper we study the first ten continuum images of MALS in depth in order to assess their quality. We investigate the systematics




Figure 3.1: Sky distribution of the first 391 observed pointings of MALS in equatorial coordinates. The Galactic plane is largely avoided, and since 89% of the pointings are selected directly from NVSS, the vast majority of pointings are above a declination of -40 degrees. The pointings used in this analysis are highlighted in red.

in calibration, imaging, and source finding on image quality and source counts, and extrapolate our findings to the rest of the survey.

This paper is organised as follows. In Section 3.2 we describe the MALS data. The initial creation of the source catalogues and the completeness measures are described in Section 3.3. In Section 3.4 we describe results from the full catalogue of sources. We investigate how different source populations affect the catalogues in Section 3.5. In Section 3.6 we assess the prospects for a dipole measurement with MALS using the results in this paper. Finally, in Section 3.7, we summarise the findings of this paper.

3.2 MALS data

The distribution of the first 391 observed pointings of MALS is shown in Figure 3.1. In order to assess the data quality of the individual MALS pointings and the impact for the dipole estimates an initial set of ten pointings, shown in Figure 3.1 in red, has been selected out of five observing runs to probe different ranges of right ascension, declination, and central source flux density.

3.2.1 Observations and calibration

The general setup of a single MALS observation includes observations of three science targets and corresponding calibrators. The observation is scheduled with a flux calibrator observed for 10 minutes at the start and end of each observing run. Each target is observed for 20 minutes at a time, cycling through all targets three times for a total observing time of an hour per source. Before and after

Table 3.1: Calibration details of the pointings presented in this paper, grouped by observation runs. The flux densities of the calibrators are reported by the CASA `fluxscale` task, which changes reference frequency based on whether the source is a flux calibrator or gain calibrator.

Flux cal	Flux density 870 MHz (Jy)	Target	Gain cal	Flux density 1365 MHz (Jy)	Reference flux ^a 1400 MHz (Jy)	Spectral index ^a	Distance from target (degrees)
J1331+3030 J1939-6342	18.829 14.095	J2023-3655	J2052-3640	1.45 ± 0.005	1.367 ± 0.017	-1.258 ± 0.02	5.8
J0408-6545 J1939-6342	27.027 14.095	J0126+1420	J0108+0134	3.24 ± 0.01	3.113 ± 0.070	-0.273 ± 0.02	13.5
J1331+3030 J1939-6342	18.829 14.095	J1133+0015 J1232-0224 J1312-2026	J1150-0023 J1256-0547 J1311-2216	2.86 ± 0.005 10.66 ± 0.05 5.5 ± 0.01	2.9 ^b 9.82 ± 0.120 4.857 ± 0.060	- -0.490 ± 0.05 -1.281 ± 0.04	4.2 6.9 1.8
J0408-6545 J1939-6342	27.027 14.095	J0001-1540 J0006+1728	J2357-1125 J2253+1608	2.12 ± 0.009 15.5 ± 0.02	1.8 ^b 16.199 ± 0.198	- -0.193 ± 0.03	4.4 17.5
J0408-6545 J1939-6342	27.027 14.095	J0240+0957 J0249-0759 J0249+0440	J0238+1636 J0240-2309 J0323+0534	0.61 ± 0.002 6.15 ± 0.003 2.85 ± 0.002	0.528 ± 0.014 5.938 ± 0.131 2.766 ± 0.062	-0.246 ± 0.05 -0.154 ± 0.03 -0.920 ± 0.01	6.7 31.2 8.5

^a <https://skaafrica.atlassian.net/wiki/spaces/ESDKB/pages/1452146701/L-band+gain+calibrators> for properties of MeerKAT L band calibrators. ^b Value from the old list of calibrators, no longer publicly available.

each target observation, a nearby gain calibrator is observed for one minute. Cycling between targets like this maximises the UV-coverage with minimal increase in overhead. Observing multiple targets in a single run is not only convenient in terms of processing, but is also critical in taking stock of systematic effects, such as flux density scale or phase errors, potentially introduced during observation or calibration. All observations have a correlator integration time of 8 s, with observations carried out in 32K mode, providing 32,768 channels with a channel width of 26.123 kHz. With a frequency range of 856 – 1712 MHz, the total bandwidth is 856 MHz, with a central frequency of 1.285 GHz.

The MeerKAT data are shipped to the Inter-University Centre for Astronomy and Astrophysics (IUCAA) in India and processed by the Automated Radio Telescope Imaging Pipeline (ARTIP). The complete deployment of ARTIP in MALS is described in Gupta et al. (2021). ARTIP presents an environment where data can be processed according to user specifications and is based on the Common Astronomy Software Applications (CASA) tasks (The CASA Team et al., 2022). Each dataset undergoes a round of basic flagging, removing known radio frequency interference (RFI) frequencies. This is followed by flux calibration, bandpass calibration, and gain calibration, each step having the possibility of additional automated flagging. The final target visibilities used for the imaging process are produced by applying the flags and calibration solutions.

As part of the overall evaluation of the individual pointings, all the available information was assessed automatically with an evaluation scheme that has been developed to trace errors of the calibration process by searching through the logging information of ARTIP. This scheme also extracts relevant information from the logs, such as the flux densities of the calibrator sources. An overview of the targets and calibrators of the ten selected pointings, organised by observation block, is shown in Table 3.1. For the gain calibrators, both the flux density determined during calibration and from a reference catalogue is listed.

3.2.2 Self-calibration and continuum imaging

For the purposes of continuum imaging, the data are averaged over 32 channels and divided into 16 spectral windows (SPWs), resulting in 64 channels per SPW. Once again, frequencies with known strong RFI are flagged (see also Figure 2 of Gupta et al., 2021). The resulting dataset has a total of 960 channels, a bandwidth of 802.5 MHz (869.3 – 1671.8 MHz), and a central frequency of 1.27 GHz as a result of the edges of the band being flagged.

As each field contains a strong point source at its centre, both phase and amplitude self-calibration can be performed (Cornwell and Fomalont, 1989). In total, three phase and one amplitude calibration steps are performed, with imaging each step to improve the local sky model. As is common with self-calibration in CASA, we use the clean components created in `tclean` as the local sky model for calibration. We iterate on the model by creating masks for `tclean` using the Python Blob Detection and Source Finder (PyBDSF; Mohan and Rafferty, 2015). Starting with a mask containing only the central source, after a set number of iterations PyBDSF is used on the image to create the mask for the full field, initially with a high S/N threshold and lowering the threshold for subsequent runs to gradually expand the model. Creating the clean masks in such a way ensures that cleaning is mostly limited to real emission, while also speeding up the imaging by limiting the cleaning area.

Though the self-calibration can be a significant improvement on the image it can also be potentially unstable. To monitor the stability of solutions, a diagnostic tool for self-calibration produces a report on the variation of relevant statistics such as noise and central source flux density in different steps of calibration. As with calibration, the logs were evaluated for errors and warnings during the self-calibration process and information relevant to assessing the quality of the products, such as percentage of flagged data and theoretical noise limit, were extracted.

Imaging is performed using Multi-term Multi-Frequency Synthesis (MTMFS; Rau and Cornwell, 2011) deconvolution with four pixel scales (0, 2, 3, and 5 pixels) to model extended emission in the images and two Taylor terms to account for the spectral shape of the sources. This produces two Taylor term images, which describe the spectral shape of the emission to zeroth and first order, respectively. As such, the zeroth order Taylor term I_0 represents the continuum flux density of the field at the reference frequency of 1.27 GHz, while the first order Taylor term I_1 describes the spectral index,

$$I_0 = I_{\nu_0}^{sky}; I_1 = \alpha I_{\nu_0}^{sky}. \quad (3.1)$$

To maintain a balance between sensitivity and resolution in the images, visibilities are weighted using Briggs weighting (Briggs, 1995) with robust value of 0. Because we are imaging with a large field of view, we use W-projection (Cornwell et al., 2005) with 128 projection planes to correct for the fact that our baselines are non-coplanar. The final data products consist of the restored, model, residual, sum-of-weights, and point spread function (PSF) images for both Taylor terms. Furthermore, spectral index, spectral index error, and mask images are also produced. The continuum images have a pixel size of $2''$ and a size of 6000 x 6000 pixels. This results in a square image of 3.3 degrees on a side. Though individual pointings have different beams, as detailed in Table 3.2, they are on average aligned in the north-south direction, with a mean major axis of $9.3''$ and mean minor axis of $6.5''$.

3.2.3 Spectral index images

The L band of MeerKAT has a bandwidth of 802.5 MHz, which is large enough to be sensitive to the spectral shape of the radio emission within the band. If this is not taken into account when imaging

the full band, this incurs a large uncertainty in flux density. The general solution for this is MTMFS deconvolution, which models the frequency dependence of the emission with a Taylor expansion. In our case, as mentioned in Section 3.2.2, we model the frequency dependence of the emission in the pointings to first order in ν . With this we can create maps describing the spectral index α , defined by the relation between flux density S and frequency $S \propto \nu^{\alpha}$ ⁷, of the emission in the image.

Although MTMFS imaging also produces a spectral index image, pixels below 5 times the peak residual are masked in this image. To retain flexibility, we therefore chose to produce the spectral index images from the Taylor term images ourselves. From the definition of the Taylor term images in Eq. 3.1, a spectral index image can be obtained using $\alpha = I_1/I_0$, from which we will be able to measure the spectral indices of sources. To keep values in the spectral index image from diverging, pixels are masked where values in the Stokes I image are below $10 \mu\text{Jy beam}^{-1}$. When measuring the spectral index in some region of the image, usually defined by the extent of a source, we assign a spectral index as the intensity weighted mean of the measured pixels in the spectral index image, with intensity weighted standard deviation as the error,

$$\bar{\alpha} = \frac{\sum_{i=1}^n I_{0,i} \alpha_i}{\sum_{i=1}^n I_{0,i}}, \quad (3.2)$$

$$\sigma_{\alpha} = \sqrt{\frac{\sum_{i=1}^n I_{0,i} (\alpha_i - \bar{\alpha})^2}{\frac{n-1}{n} \sum_{i=1}^n I_{0,i}}}. \quad (3.3)$$

If more than half of the measured pixels in a region are invalid in the spectral index image, this carries over to the measured spectral index and uncertainty by assigning a masked value.

3.2.4 Primary beam correction

Due to the primary beam response of the MeerKAT antennas, sources away from the pointing centre appear fainter than they are in reality. As this effect is not corrected for in the imaging stage, resulting continuum images will have accurate flux densities at the pointing centre but attenuated flux densities that become fainter the farther from the pointing centre they are located. A simplified model of the primary beam is described in Mauch et al. (2020), which assumes the primary beam of MeerKAT as directionally symmetric, describing it with a cosine-tapered illumination function,

$$P(\rho, \theta_{pb}) = \left[\frac{\cos(1.189\pi\rho/\theta_{pb})}{1 - 4(1.189\rho/\theta_{pb})^2} \right]^2. \quad (3.4)$$

Here ρ is the distance from the pointing centre and θ_{pb} represents the angular size of the FWHM of the primary beam, a quantity that is dependent on the observing frequency, ν ,

$$\theta_{pb}(\nu) = 57'.5 \left(\frac{\nu}{1.5\text{GHz}} \right)^{-1}. \quad (3.5)$$

At the central frequency of our continuum images of 1.27 GHz, the FWHM of the primary beam is $\theta_{pb} = 67'$. This simplified model is implemented in the `katbeam`⁸ PYTHON package. As the primary

⁷ Note that throughout this chapter we use the inverse definition of α from Equation 2.1

⁸ <https://github.com/ska-sa/katbeam>

Table 3.2: Details on all ten pointings after complete processing and source finding.

Target	RA	Dec	Flux density NVSS (mJy)	Flux density MALS (mJy)	Spectral index	PSF maj "	PSF min "	PSF PA °	Counts	σ_{20} ($\mu\text{Jy beam}^{-1}$)	Demerit (mJy)
J0001-1540	00 ^h 01 ^m 41 ^s .57	-15°40'40".60	436	513.6 ± 1.0	-0.85 ± 0.04	7.7	6.3	-8.9	2132	26	14.5
J0006+1728	00 ^h 06 ^m 47 ^s .35	+17°28'15".40	226	220.1 ± 0.7	-0.28 ± 0.08	11.4	6.3	-7.1	1378	29	10.1
J0126+1420	01 ^h 26 ^m 13 ^s .24	+14°20'13".10	577	685.6 ± 0.7	-0.95 ± 0.09	10.5	6.3	-2.5	1591	33	8.6
J0240+0957	02 ^h 40 ^m 27 ^s .19	+09°57'13".00	521	589.3 ± 0.7	-1.11 ± 0.09	10.2	6.6	9.1	986	48	18.7
J0249-0759	02 ^h 49 ^m 35 ^s .41	-07°59'21".00	646	711.0 ± 0.6	-0.97 ± 0.09	9.2	6.6	-1	2619	19	7.7
J0249+0440	02 ^h 49 ^m 39 ^s .93	+04°40'28".90	420	472.6 ± 0.3	-0.80 ± 0.09	8.1	6.7	-7.6	1558	29	6.3
J1133+0015	11 ^h 33 ^m 03 ^s .12	+00°15'48".90	233	377.7 ± 0.9	-0.01 ± 0.07	8.9	6.7	-15.3	803	52	15.9
J1232-0224	12 ^h 32 ^m 00 ^s .13	-02°24'04".10	1647	1823.4 ± 5.2	-0.31 ± 0.09	8.5	6.7	-9	611	73	20.8
J1312-2026	13 ^h 12 ^m 07 ^s .86	-20°26'52".40	778	851.4 ± 0.4	-0.83 ± 0.09	7.7	6.3	-9.1	2431	21	12.2
J2023-3655	20 ^h 23 ^m 46 ^s .21	-36°55'21".20	436	406.6 ± 0.5	-0.11 ± 0.12	10.9	6.8	-88	2160	22	5.8

beam is frequency dependent, it affects the spectral index images, increasing the measured spectral index away from the pointing centre. The spectral index change induced by the primary beam can be approximated by

$$P_\alpha(\rho, \nu) = -8 \log(2) \left(\frac{\rho}{\theta_{pb}} \right)^2 \left(\frac{\nu}{\nu_0} \right)^2, \quad (3.6)$$

Again, we assume the frequency ν to be equal to the central frequency $\nu_0 = 1.27$ GHz.

In reality, the MeerKAT primary beam in the L band is more complicated and cannot be completely described by a directionally symmetric model. Villiers and Cotton (2022) present and analyse holographic measurements of the MeerKAT primary beam, showing the directional asymmetries present due to variations between individual antennas. For an accurate model of the primary beam, we use these holographic measurements to correct our images. As we utilise the full 802.5 MHz bandwidth of the L band for these images, a primary beam correction must take this into account. Though a wideband primary beam correction is implemented in the CASA task `widebandpbcor`, there are no models of the MeerKAT beams available. As such, we implement the wideband primary beam correction ourselves using the same basic recipe, which consists of creating a primary beam with a frequency structure matching that of the image, in this case creating a primary beam model for each of the 16 SPWs of the continuum data. As in the imaging step, we model the multi-frequency primary beam with two Taylor terms. The primary beam corrected Taylor term images are then defined as follows:

$$I'_0 = P_0^{-1} I_0, \quad (3.7a)$$

$$I'_1 = P_0^{-1} \left(I_1 - \frac{P_1}{P_0} I_0 \right). \quad (3.7b)$$

Here, P_0 and P_1 represent the zeroth and first order Taylor term primary beams, respectively, where P_0/P_1 should be equal to α_{pb} as specified in Eq. 3.6.

While we use the holographic wideband primary beam corrections described in Eq. 3.7 for the main results of this work, we also briefly explore the simplified corrections of Eqs. 3.4 and 3.6 and see how they compare to the wideband corrections. At applying the primary beam corrections, the image is cut off at the 5% level of the primary beam (at the central frequency of 1.27 GHz), which leaves us with a circular image with a diameter of approximately 4000 pixels, or 2.2 degrees. As a result of reduced sensitivity towards the edges of the image, the noise is increased there.




Figure 3.2: Flux offsets of the central sources and gain calibrators per field. Gain calibrators are compared to their reference flux density (circles) as specified in Table 3.1 and with their NVSS counterparts (diamonds). Central sources are compared to their NVSS counterparts (squares), and NVSS flux densities are all converted to rest frequency assuming $\alpha = -0.75$. Sources are ordered by observing run, separated by vertical black lines.

3.2.5 Assessment of calibration

Processing the raw data to the final scientific data products can introduce errors, affecting the flux density scale. A first order estimation of the flux density scale can be obtained by comparing the flux densities of the gain calibrators with their literature values. As discussed in Section 3.2.1, we evaluate the automated process of the data calibration by generating diagnostic reports and automatically evaluating logged information in order to determine problems in the data processing. This evaluation singles out errors and warnings present in the logs, allowing direct insight into any problems that might have occurred during the calibration process. Furthermore, it extracts information we can use to assess the quality of calibration from the logs, such as the flux densities of calibrator sources.

Table 3.1 summarises the observation and calibration details, showing the targets and their associated calibrator sources. The flux densities of the flux- and gain calibrators are extracted from the logs and the flux densities of gain calibrators are compared to the MeerKAT reference catalogue (Taylor and Legodi, 2021). We extend this to a broader assessment of the flux density scale in Figure 3.2, where we show the flux density offsets of the gain calibrators and central sources of the individual pointings. Along with the comparison in Table 3.1, both gain calibrators and central sources (see Table 3.2) are compared to their NVSS counterparts. Flux densities are corrected for frequency using the spectral index from the reference catalogue if available, assuming $\alpha = -0.75$ otherwise. Combining the measurements from the ten pointings, the mean flux density ratios are 1.03 ± 0.26 between the gain calibrators and their NVSS counterparts, 1.07 ± 0.07 between gain calibrators and their reference values, and 1.08 ± 0.19 between central sources and their NVSS counterparts. We note




Figure 3.3: Median stacked pixel values of the rms noise images of all ten pointings. As primary beam correction is applied, the noise goes up towards the edges of the image. Since a strong central source is always present, the noise is always higher in the centre as well. The given contours from white to black are 20, 40, and 60% rms noise coverage. The stacked beam (50 time increase in size) of all the pointings is shown in the lower-left corner, matching the elongated structure in the centre.

that the absolute amplitude calibration of NVSS is based on Baars et al. (1977) and has an uncertainty of up to 12% with respect to the here used Perley and Butler (2017) scale, depending on the calibrator used.

We note that the SUMSS and NVSS measurements were taken with different instruments at different times, so some variation is to be expected. The current assessment does not include astrometric precision, as calibrators are not imaged. We assess this aspect along with the another flux density scale assessment by cross-matching the full catalogue of sources with other surveys in Section 3.4.2.

3.2.6 Assessment of image quality

With any radio image, there is a great number of variables that can influence the quality of the image, both related to intrinsic properties of the pointing and to the process of calibration and self-calibration. As discussed in Section 3.2.2, a report is generated that monitors image statistics such as noise and central source flux during different self-calibration steps. Furthermore, the logs are automatically evaluated for possible errors and warnings and information relevant to the quality of the self-calibration and imaging is extracted. To evaluate the final image product, the image quality of the individual ten pointings is assessed by using the root mean square (rms) noise maps that are automatically produced during the source finding procedure by PyBDSF (see Section 3.3). In particular, we investigate the overall noise characteristics by evaluating the sky coverage with respect to the rms noise. A direct




Figure 3.4: Rms noise coverage for all ten pointings. The dotted line indicates the 20% coverage level, which is used to define σ_{20} . Noise varies appreciably between the pointings; however, the overall structure of the rms noise coverage curves remains consistent, indicating that σ_{20} is a good zeroth order measure of the noise scale.

measurement of the noise allows us then to easily correlate image quality with other characteristics of the pointings.

We create a smoothed representation of the ten pointings by median stacking their normalised rms noise images, which is shown in Figure 3.3. As all pointings have a strong source at their centre, the noise is increased at the pointing centre and increases towards the edges of the pointing as a consequence of the primary beam response. Figure 3.3 shows that some directional effects are left in the image. Notably, there is an elongated noise structure in the centre, associated with the bright central source, aligned in the north-south direction. The stacked beam included in the figure aligns well with the elongated structure, indicating that the most prominent structures are a result of the shape of the stacked PSF of the images. The imprint of the stacked PSF is also the most likely cause of the cross-like structure seen in the stacked image. Though we have the wideband primary beam correction based on holographic images that take into account the asymmetries present in the primary beam, pointings are observed for three separate blocks of 20 minutes in an observing night, which smears out the asymmetries in the primary beam⁹. This effect cannot be easily corrected for in the image plane, but could be taken into account during imaging using A-projection (Bhatnagar et al., 2008). Though present, the asymmetries here are small and dominated by the other noise structures in the image.

The usual method of determining rms noise in an image relies on measuring rms noise in an area close enough to the pointing centre to not be affected by the primary beam and far enough from strong sources to be unaffected by artefacts. Due to the number and structure of MALS pointings, this cannot be reliably done in an automated fashion. Instead, we investigate the differences in noise

⁹ MeerKAT antennas have Alt-Az mounts, such that the sky rotates with respect to the dish while observing

level between individual pointings by using the rms noise images to assess the rms noise coverage, measuring the cumulative distribution of noise levels across an image. Figure 3.4 shows the rms noise coverage curves for the individual fields, and it can be seen that rms noise coverage curves are similar in structure but offset from each other. To quantify this offset and thus characterise the noise in the individual pointings, we define σ_{20} at 20% rms noise coverage, representing the noise in the central portion of the image (see Figure 3.3). We will see that σ_{20} excellently serves as a normalisation factor to account for the differences in noise levels between the pointings, and can be used to unify the assessment of individual pointings and extend them to the full survey.

There are several factors which can contribute to the overall noise level in an image, not all of which are easily quantifiable. However, an important aspect to consider is the shape of the synthesised beam or PSF, determined by the UV-coverage of the observation, which in turn is determined by the array configuration, observing time, and elevation of the target at the time of observation. There are two aspects to the PSF that influence image noise. A measurement of rms noise in Jy/beam will be influenced by the shape of the beam¹⁰, and very bright sources can have persistent and bright sidelobes from the shape of the PSF that are difficult to clean completely and as a result push up the noise in an image. To quantify this last effect we calculate the demerit score detailed in Mauch et al. (2020) to estimate the contributions of bright sources to rms noise in the image. We calculated the independent source contributions to the errors in the image using all sources that have an unattenuated flux density of more than 100 mJy. The demerit score, d , is then defined as

$$d = \left[\sum_{S > 100 \text{ mJy}}^i \left\{ \left(\frac{8 \ln(2) \rho \sigma_p}{\theta_{pb}^2} + \sigma_g \right) S_{a,i} \right\}^2 \right]^{1/2}, \quad (3.8)$$

where the first term represents the contribution of pointing error σ_p scaling with distance from the pointing centre ρ , and the second term is the receiver gain error σ_g . The contribution of each source comes in the form of their attenuated flux density S_a . Appropriate values for the MeerKAT L band are detailed in Mauch et al. (2020), which we also use ($\theta_{pb} = 67'$, $\sigma_p = 30''$, $\sigma_g = 0.01$). The demerit scores of all pointings are included in Table 3.2. A correlation is present between demerit score and σ_{20} , and especially pointings with high σ_{20} show increased demerit scores. Though pointings with lower σ_{20} show more scatter in their demerit scores, this nonetheless shows demerit score as a first order estimate of pointing quality, which we can utilise as a predictive measure.

3.3 Source finding

With thousands of sources expected to be detected in every MALS pointing, we require an automated source finding algorithm to find and characterise these sources. A small number of these are suitable for radio images, and perform comparatively similar (Hale et al., 2019). Of these, PyBDSF has been used in several recent data releases of large-scale surveys, such as the LOFAR Two-Metre Sky Survey (LoTSS; Shimwell et al., 2019) and the RACS (Hale et al., 2021). PyBDSF stands out in its ability to model extended emission with its wavelet decomposition module, and provides easy ways to compile source catalogues and assess the quality of the fields. Besides generating catalogues, PyBDSF provides output maps related to the input image, such as the rms noise images we used in Section 3.2.6,

¹⁰ The clean beam of an image is determined during imaging by fitting a 2D Gaussian to the central lobe of the PSF.

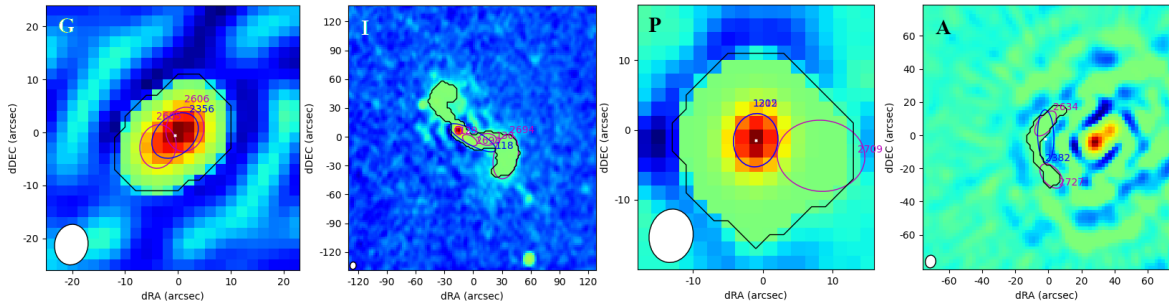


Figure 3.5: Examples of possible source classes. The black outline shows the island threshold, the magenta ellipses show the individual Gaussian components fit to the source, and the blue ellipses show the combined Gaussian describing the source. From left to right: (i) An elongated source fit by two Gaussian components. The combined Gaussian describes the source adequately, and it has been assigned the ‘G’ class. (ii) A likely FRI source with complex structure, better described by the island than the Gaussian components. It has been assigned the ‘I’ class. (iii) A point source with an additional noise peak that has been fitted with a Gaussian component. It has been assigned the ‘P’ class. (iv) An artefact caused by a nearby bright source. It has been assigned the ‘A’ class.

and mean and residual images. Once rms noise and mean maps are obtained PyBDSF allows these maps to be used as input to ensure source finding is performed with the exact same parameters. For MALS, we thus make use of PyBDSF, both for creating clean masks during self-calibration as detailed in Section 3.2.2, and integrating PyBDSF into the workflow to automatically carry out source finding, cataloguing, cross-matching and combining catalogues, using PYTHON-based scripts developed by the authors¹¹.

In order to understand the impact of the individual pointings to a general catalogue, we evaluate the source finding procedure for each pointing. We investigate completeness (what fraction of sources do we detect) and purity (what fraction of sources is real) in source counts with respect to signal-to-noise ratio, flux density, and source size, as well as PyBDSF’s capability to accurately recover flux densities.

3.3.1 Stokes I catalogues

In order to compile a source catalogue from PyBDSF, various steps that depend on the initial setup of PyBDSF are needed. PyBDSF identifies islands of emission that are brighter than the island threshold. Within these islands PyBDSF finds emission peaks above a corresponding pixel threshold, and for each peak found fits a 2D Gaussian to the peak and surrounding emission. Performing source finding on our MALS images, we impose an island threshold of 3σ , and a pixel threshold of 5σ . Individual Gaussian components are combined into sources in a way that can be specified by the user, and we elect to combine Gaussian components that occupy the same island into a single source. The rms noise in the images is determined by a sliding box, and we decrease the size of the sliding rms box near bright sources to avoid spurious detections of artefacts around these sources as much as possible. Furthermore, to improve fitting of extended sources in the field, we enable à trous wavelet decomposition (Holschneider et al., 1989). The PyBDSF settings can be summarised as follows:

```
thresh_isl = 3.0
```

¹¹ <https://github.com/JonahDW/Image-processing>




Figure 3.6: Assigned flags to sources with multiple Gaussians, separated by the number of Gaussian components fit to them by PyBDSF. Keeping in mind that most of these sources are fit with two Gaussian components, the ‘I’ class is preferred for sources with three or more components, while the ‘P’ class consists almost exclusively of sources with two components.

```

thresh_pix = 5.0
rms_box = (150,30)
adaptive_rms_box = True
adaptive_thresh = 100
rms_box_bright = (40,15)
group_by_isl = True
atrous_do = True
atrous_orig_isl = True
atrous_jmax = 3

```

For the purposes of analysing and building the final catalogue, we required a number of output products from PyBDSF. The output from source finding includes both a catalogue of sources and of individual Gaussian components. Furthermore, a background rms noise and background mean image are produced, as well as a residual image.

For a single pointing, we ran source finding using PyBDSF and modified the output source catalogues by adding to the existing columns the ID of the MALS pointing, a source name following IAU convention, and the distance of the source to the pointing centre. As PyBDSF does not calculate spectral indices unless an image has multiple channels, we measured the spectral index of the sources in our fields from the spectral index images as described in Section 3.2.3, using the extent of the Gaussian (major axis, minor axis, position angle) of the source to define the region in the image.

Though PyBDSF is configured to avoid spurious detections as much as possible, it is unavoidable

that some artefacts are falsely identified as sources. We identified artefacts around the ten brightest sources in each image by flagging sources within five times the major axis of the beam that have less than 10% of the peak flux density of the bright source. This is largely motivated by the shape of the PSF, which can have sidelobes with a strength of up to 10% of the maximum. Though this does not get rid of all false detections in the image (see Section 3.3.2), it flags the most prominent imaging artefacts.

To assess the quality of the Gaussian fitting by PyBDSF, we performed a visual inspection on select sources. We created cutouts from the images and performed visual inspection, which was implemented in a separate module based on PYTHON and CASA¹². PyBDSF assigns each source a flag indicating whether the source is fit by a single Gaussian ('S'), multiple Gaussian components ('M'), or Gaussian component(s) on an island with other sources ('C'). Since all Gaussian components that occupy the same island are always combined into one source, the 'C' flag is not present in our catalogues. For the visual inspection, we considered all sources made up of multiple Gaussian components. As such, all sources that carry the 'M' flag — which make up 8% of the all sources found in the fields — are flagged for visual inspection. Through the visual inspection, we then assigned an additional flag indicating the nature of the source and how well it is described by the PyBDSF model:

- G: Sources that are well described by the Gaussian model.
- I: Complex sources that are not adequately described by the Gaussian components fit to them. The flux density of these sources is better described by the integrated flux of the island, and their position by the flux weighted mean position of the island.
- P: Sources fit with multiple Gaussian components where only one is required to adequately describe the source. Other Gaussian components are likely fit to noise fluctuations coinciding with the source.
- A: Artefacts that will be flagged as such in the catalogue.

Figure 3.5 shows an example for each of the cutout classes, and how we identify the different possible cases. To aid in visual inspection, in the source cutout we plot the individual Gaussian components, the combined source Gaussian, and the island threshold. Therefore, in this step we use both the source catalogues and the Gaussian component catalogues. Along with the cutout classes, additional columns are added to the table that describe the sources. In the cutouts we measure the integrated flux density of the island, the spectral index of the island, the intensity weighted mean position of the source, and a flag indicating if these measures are valid. Additionally, the number of Gaussian components is recorded for each source, as the initial PyBDSF catalogue only indicates whether a source has been fit with multiple Gaussian components or not.

Figure 3.6 shows the classification of all sources in the ten pointings that have been fit with multiple Gaussian components, 1259 in total. We see that almost all 120 sources assigned with the 'P' class have two Gaussian components, and a relatively large percentage of sources with the 'I' class have more than two Gaussian components assigned. Around 185 (15%) of these sources were considered to be adequately described by their Gaussian components, while 946 (75%) are more complex and better described by their island attributes. Only 8 sources are flagged as obvious artefacts.

¹² <https://github.com/JonahDW/CASA-Poststamp>

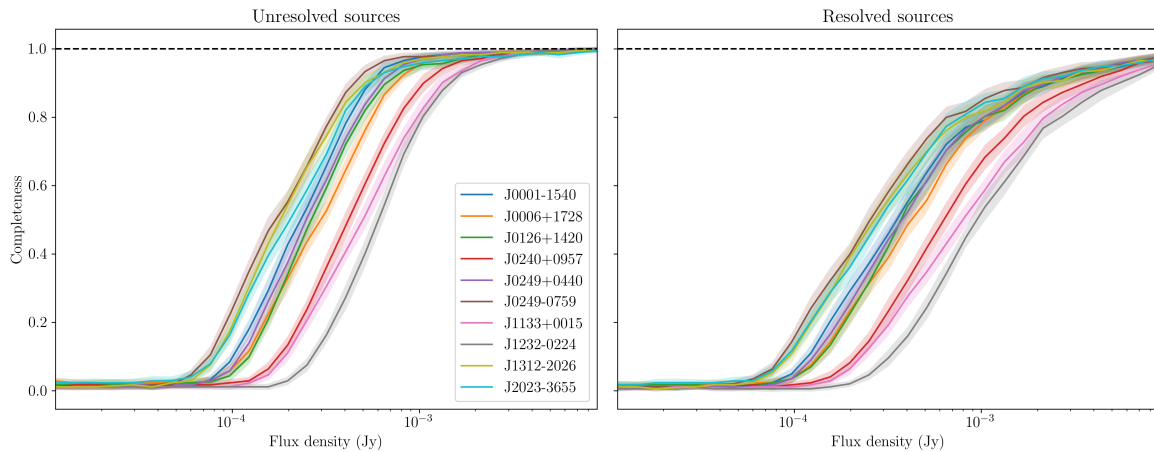


Figure 3.7: Completeness for unresolved sources (left) and resolved sources (right) for the different fields and their associated uncertainties as a function of flux density. There are large differences between the pointings: pointings with higher noise levels have lower overall completeness. The completeness is lower for resolved sources as well.

3.3.2 Evaluation of individual pointings

In order to determine the reliability of the source finding routine and to assess how detection of sources is affected by their properties, we measured the completeness, purity, and flux recovery statistics of the catalogues. Here we assess these qualities for individual pointings to see how characteristics of the pointings such as central source flux density and noise level affect these quantities.

To measure completeness and flux recovery, we required complete knowledge of the intrinsic flux densities and shapes of the sources that are present in the image. To that end, we used realistic samples of simulated extragalactic radio sources from the Wilman et al. (2008) simulation of the SKA Design Study (SKADS). Though more recent simulations such as the Tiered Radio Extragalactic Continuum Simulation (T-RECS; Bonaldi et al., 2019) are available, the SKADS catalogues include morphology details of all sources and source components, which is necessary information when injecting sources into the data. From the SKADS simulations we created mock catalogues with 5000 sources uniformly distributed in flux density that have a flux density above $10 \mu\text{Jy}$, which is equal to the limit of thermal noise ($10 \mu\text{Jy}/\text{beam}$) for an unresolved source. With this we allowed for the possibility of noise fluctuations to push sources above the detection threshold. We injected sources from the mock catalogue uniformly distributed into the residual images produced by PyBDSF, which are devoid of sources but share the noise characteristics of the original images. We then performed the source finding routine again on these images, using the same mean and rms noise maps determined by PyBDSF from the original image. This ensures that source finding is performed in the exact same way as the original image. We considered a source recovered if it is detected within the FWHM of the major axis of the clean beam from the original position. In order to reach a more robust measure, this process was repeated 50 times for each pointing, separately for point sources and resolved sources.




Figure 3.8: Completeness for unresolved sources as a function of flux density and distance from the pointing centre, ρ , for the fields J0249-0759 (left) and J1232-0224 (right). The radially averaged 5σ curves (black lines) for the corresponding pointings are seen to follow the zone where completeness transitions from zero to one. Due to the presence of a strong central source in J1232-0224, completeness is lower in the central region of this pointing. Pixels with no sources in them have been coloured grey.

Completeness

With the procedure described above, we can make a statistically robust assessment of the completeness in the pointings. The (source) completeness in this case simply gives the fraction of sources that is detected, most commonly measured as a function of flux density of the source. The completeness curves for the individual pointings, for both resolved and unresolved sources, can be appreciated in Figure 3.7. Not only is there a large difference between resolved sources and unresolved sources, pointings individually have large differences between them as well. To investigate other aspects of the completeness, we look at the fields J0249-0759 and J1232-0224, which have the lowest and highest noise levels among the pointings, respectively (see Table 3.2), which should yield the most extreme cases and allow us to probe variation between the fields.

Unresolved sources: As the SKADS catalogues describe the intrinsic shapes of sources, we can assess completeness for point sources by only injecting sources with a major axis of zero. The sources are defined in the image as delta functions, and convolved with the clean beam of the individual image. As the total flux density of point sources is concentrated in one peak, they are much easier to detect relative to resolved sources. Point sources allow us to assess completeness without being affected by source morphology, and so we use them to determine the completeness with respect to distance from the pointing centre. As sensitivity decreases outwards from the centre, we expect completeness to decrease as well. Figure 3.8 shows source completeness as a function of both flux density and distance from the pointing centre for the pointings J0249-0759 and J1232-0224. It is clear that indeed the completeness decreases with increased distance from the pointing centre, but is also lower near the pointing centre. This is a direct result of the strong source at the centre of each pointing pushing up the noise in its immediate vicinity. In the case of J1232-0224, which has a very strong source at the pointing centre, there is significant impact on completeness in the central portion of the image. To investigate the relation between the completeness and noise floor as a function of distance to the




Figure 3.9: Completeness for resolved sources as a function of flux density and ratio between the area of the source and the beam Q_A , for the fields J0249-0759 (left) and J1232-0224 (right). Completeness can be seen to linearly decrease in the log-log scale as a function of Q_A , showing that larger sources are harder to detect. Flux densities are compensated for the local noise in order to equalise completeness for different positions in the image. Pixels with no sources in them have been coloured grey.

pointing centre we use the rms maps created by `PyBDSF`. By radially averaging the rms noise of the pointing, we obtain rms noise as a function of distance to the pointing centre. As we have set the detection threshold at 5σ , we plot the radially averaged 5σ detection curve in Figure 3.8, showing that this curve almost perfectly follows the completeness ‘transition zone’ for both pointings. In this transition zone the completeness goes up steeply from zero to one, and the flux density at which this occurs is directly related to the noise floor.

Resolved sources: We performed the same experiment for resolved sources, where we define a resolved source as a source that has major axis and minor axis larger than 0 in the SKADS catalogues¹³. Sources are randomly selected out of the catalogue, so the distribution of source shapes injected in the image represents the distribution of the SKADS sample. These sources are injected as Gaussians into the image, and as with point sources, convolved with the clean beam. Owing to their lower surface brightness, resolved sources are often less easily detectable compared to point sources with the same flux density. To check how the size of sources affects completeness, we define the area ratio Q_A of a source as the ratio between the area of the source and the beam as defined by their Gaussian characteristics. These are the major and minor axes θ_{maj} and θ_{min} for the source and B_{maj} and B_{min} for the clean beam of the pointing,

$$Q_A = \frac{\theta_{maj}\theta_{min}}{B_{maj}B_{min}}. \quad (3.9)$$

We show the completeness as a function of area ratio and flux density in Figure 3.9 for the pointings J0249-0759 and J1232-0224. As shown in Figure 3.8, the completeness of unresolved sources is

¹³ A small subset of sources with minor axis of 0 and major axis larger than zero are not included in the simulations.




Figure 3.10: Input flux density plotted against the measured flux density, for the fields J0249-0759 (left) and J1232-0224 (right), based on 100 simulations. As the noise floor in J1232-0224 is relatively high ($\sigma_{20} = 80 \mu\text{Jy}/\text{beam}$), only sources above $200 \mu\text{Jy}$ are detected, while in J0249-0759 sources are detected down to $50 \mu\text{Jy}$.

related to the noise floor. This same relation should be present for resolved sources, in addition to the relation between completeness and source size. In order to disentangle the two different contributions to completeness for resolved sources, we divided flux densities by the ratio of local noise to the lowest noise in the image. The result in Figure 3.9 shows completeness for uniform noise, so that only the source size and flux density affect completeness. We see a power law decrease (linear in the log-log scale) in completeness as a function of area ratio in both pointings, indicating that this is a universal feature for our source detection. This can be easily understood by considering that for larger sources the total flux density is divided over a larger area, which decreases the peak flux density that is used to detect these sources.

Flux recovery

Using the mock catalogues, we investigate the ability of PYBDSF to accurately recover flux densities. This can be checked by looking at the flux densities measured by PYBDSF relative to the input flux densities from the mock catalogues. This is an important quality to verify as deviations from the expected 1:1 relationship are obviously undesirable. In Figure 3.10 we show the measured flux densities against input flux densities of the pointings J0249-0759 and J1232-0224. We see that on average sources have a flux density that matches with their input value. There is however a portion of sources with lower input flux density that have a significantly higher measured flux density than their input. These sources have their flux densities boosted by noise fluctuations, which are present in various orders of strength in the images, from thermal noise to calibration artefacts. We expect these sources to land on positive as well as negative noise peaks, but only sources on positive peaks will be detected. This results in an Eddington bias (Eddington, 1913) pushing up the distribution of flux densities. To make quantitative statements about this bias, we need to combine data from all the 10 pointings, which we do in Section 3.4.3.

Limitations of simulations

The method we have used here for measuring completeness and flux recovery relies on injecting sources directly into the residual images and measuring their properties with P_YBDSF. The advantage of this is a direct probe into the machinery of P_YBDSF, as this is the only ‘black box’ between the input sources and the measurement. However, these simulations ignore some effects that affect the flux densities and shapes of sources in radio data, such as calibration effects, clean bias, and averaging effects like time and bandwidth smearing (Bridle and Schwab, 1999). Probing these requires injecting sources directly into the visibilities and reprocessing the image, something that is not efficient for a large survey such as MALS. Finally, sources are injected into the image convolved with the clean beam as opposed to the PSF. This is well motivated for brighter sources, as these have been mostly cleaned during the imaging process. For faint sources this is not the case, especially since the masks for cleaning are generated by P_YBDSF and are thus subject to the same selection that we used for the final images. To make the simulations more realistic, all undetected sources should therefore be convolved with the PSF. It is not clear how this should affect source finding, but the general consequence of this is that below the detection threshold sources immediately become fainter as a consequence of being convolved with the PSF rather than the clean beam. The PSF also spreads the emission of these sources over a large area, which could affect rms noise if source crowding is high enough. This would however only be the case if images would be close to or at the confusion limit, which is not the case for MeerKAT in the L band down to at least 0.25 μ Jy (Mauch et al., 2020).

Purity

The purity, or inversely the false detection rate, measures what fraction of the sources detected in the image are true detections. For a well chosen detection threshold, the amount of false detections in an image is expected to be small. It is important to have a handle on the amount of false detections, as it should be taken into account when calculating number counts. To determine the purity, we invert the pixel values of the images and run P_YBDSF on the inverted images, using the rms and inverted mean maps determined by P_YBDSF from the original image. This again ensures that source finding is performed in the same way as on the original image. Since all real sources have positive flux density the only sources detected in the inverted images will be false detections. These false detections broadly fall into three categories, which we differentiate as noise peaks, (calibration) artefacts, and ghost sources. Noise peaks are statistical outliers of noise and can therefore appear at any point in the image, and are symmetric around the mean, such that these sources detected in the inverted image correspond roughly to the false detections in the normal image. Artefacts are sidelobes found around strong sources in the image, making them more easily traceable. As described in Section 3.3.1, we consider a source to be an artefact if they are found within 5 times the major axis of the clean beam of the ten brightest sources in the image, and have less than 10% of the peak flux density of the brightest source. As the brightest negative sidelobe of the PSF is in general twice as bright as the brightest positive sidelobe, we would expect more artefacts to be found in the negative image. This seems consistent with the data, as using this criterion for artefacts flags 44 of the 241 sources found in the inverted images, while flagging 22 sources in the pointing catalogues. Finally, there are ghost sources, which appear as negative sources too bright to be noise fluctuations, in some cases even strong enough that they have sidelobes that are detected as sources. These sources can be caused by calibration with an incomplete sky model, and only have faint positive counterparts (Grobler et al., 2014). Strong




Figure 3.11: Purity of catalogues for different pointings. The fraction of false detections is indicated by the red line. The dashed red line indicates the fraction of false detections without flagging artefacts. The open histograms show the number of sources detected in the pointings, with the filled histograms indicating the number of false detections.

ghosts can add to the number of false detections with their sidelobes, but only a handful of such cases are seen in the images.

We plot the amount of false detections per pointing in Figure 3.11, both in terms of absolute counts (coloured bars) and fraction (red line). The amount of false detections strongly depends on pointing, and we find two pointings that are most strongly affected: J0001-1540 and J1232-0224. The latter is affected by a strong central source, which leads to reduced number counts and an increased fraction of false detections, while the former would be considered a good pointing, both in terms of number counts and noise properties. There does appear however a cluster of relatively bright (10-100 mJy) sources present far from the pointing centre, which can contribute to noise. The presence of a number of strong sources far out in the field has also potentially affected self-calibration, as a high number of ghost sources are seen in the image. This result suggests that purity of any individual pointing is not always easily predictable, and should each be assessed separately.

3.4 Combined catalogue

The combined source catalogue of ten pointings contains 16,307 sources, and covers 35.7 square degrees of sky. In the previous section we mostly assess the quality of individual pointings. Here we combine the catalogues of the individual pointings to increase statistical power, which allows us to investigate subtler systematic effects that affect all pointings.

3.4.1 Correcting residual primary beam effects

In Section 3.2.4 we described primary beam corrections to both the flux densities and spectral indices in the images. Besides the main wideband primary beam corrections using holographic images, we also described corrections with simplified analytic forms. Before investigating the difference between these methods, we must make additional corrections to residual primary beam effects. In general, the simplified analytical corrections work well up to the FWHM of the primary beam, but farther out results begin to diverge. This is an effect that is seen in both the spectral indices as well as the flux densities of sources, mainly caused by using the primary beam correction based on the central frequency $\nu_0 = 1.27$ GHz for the entire bandwidth of 802.5 MHz. In order to take into account the contribution of the entire bandwidth, we recalculate the corrections by integrating over the bandwidth rather than assuming the frequency to be equal to ν_0 . The necessary corrections are computed for each source in the catalogue separately, depending on distance from the pointing centre and spectral shape.

In the images corrected by the simplified analytical function of Eq. 3.4, flux densities of sources appear higher the farther they are from the pointing centre. In order to properly correct for the effect in flux density, the spectral index of the source must be known or assumed. For reasons explored in Section 3.5.2, we cannot trust all spectral indices to be accurate and perform this correction assuming $\alpha = -0.75$ for all sources. We then calculate a correction factor for the flux densities as a function of distance from the pointing centre. The assumed primary beam model is as before (Eq. 3.4), and the correction is computed by integrating the primary beam over the frequency range of the band. Considering a source with some spectral index α , the flux density of the source is described by $S(\nu, \alpha) \propto \nu^\alpha$. Due to the effect of the primary beam, the flux density of the source has some attenuation factor $a(\rho, \alpha)$ applied to it. This factor is described by the primary beam:

$$a(\rho, \alpha) = \frac{\int_{\Delta\nu} S(\nu, \alpha) P(\rho, \nu) d\nu}{\int_{\Delta\nu} S(\nu, \alpha) P(0, \nu) d\nu}. \quad (3.10)$$

Since the flux densities have already undergone primary beam correction, we need to correct for the ratio between this term and the correction from Eq. 3.10,

$$S_{corr} = \frac{P(\rho, \nu_0)}{a(\rho, \alpha)} S_{measured}. \quad (3.11)$$

The effect of this correction should become visible when comparing flux densities to external catalogues. If the flux densities are properly corrected, the flux density ratio between catalogues should be constant as a function of distance to the pointing centre.

In contrast to the flux densities, both the analytical function from Eq. 3.6 and the wideband primary beam correction from Eq. 3.7 leave a residual effect in the spectral indices of sources farther from the pointing centre. Taking Eq. 3.6 to describe the spectral index induced by the primary beam variation, we correct these values taking the full bandwidth into account, recalculating the effect of the primary beam on spectral indices by integrating over the bandwidth, $\Delta\nu$:

$$P_{\alpha,int}(\rho) = -8 \log(2) \left(\frac{\rho}{\theta_{pb}} \right)^2 \int_{\Delta\nu} \left(\frac{\nu}{\nu_0} \right)^2 d\nu. \quad (3.12)$$

To correct the already measured spectral indices present in the catalogues we subtract the difference

between the integrated and original primary beam correction term,

$$\alpha_{corr} = \alpha_{measured} + [P_{\alpha,int}(\rho) - P_{\alpha}(\rho, \nu_0)]. \quad (3.13)$$

3.4.2 Cross-matching catalogues

To further investigate systematics that affect the pointings on a more general level, we continue the assessment from Section 3.2.5, now using the sources of the entire field. We cross check our sources with their counterparts from NVSS and RACS, as all our pointings here are within the sky coverage of these two surveys. Cross-matching was performed by checking whether source ellipses, defined by the 3σ extent of the Gaussians describing these sources, overlap between the catalogues. We required a minimum overlap in area of 80% to consider sources to be a match. Sources in one catalogue could be matched with any number of sources in the other, to account for different resolutions between the catalogues. Due to uncertainties in position and flux density near the NVSS detection threshold of 2.5 mJy, sources below 5 mJy in NVSS were not considered. We find that 997 sources are matched to NVSS, of which 845 are matched to a single source, and 2064 sources are matched to RACS, with 1949 matched to a single source.

There are a number of factors that can influence astrometric precision of an observation, such as errors in the reference frequency or timestamps. Some of these errors were present in earlier MeerKAT observations (e.g. Mauch et al., 2020; Villiers and Cotton, 2022). While the errors should no longer be an issue, it is important to cross check positions in the field with an external catalogue for potential astrometric errors. The astrometric offsets of sources to their NVSS counterparts can be seen in Figure 3.12, where the offsets are shown only for single matched sources. Overall, the offsets are very small with a median offset of $\sim 0.3''$, which is less than one-sixth of the image pixel size ($2''$) and well within the uncertainty. The scatter in both directions is smaller than $3''$, which is less than the semi-minor axis of the average clean beam of $3.25''$, also shown on the figure.

In Section 3.4.1 we corrected spectral indices and flux densities accounting for residual effects introduced by the frequency range covered by the band. Cross-matching sources with external catalogues is an important check of the correctness of their measured flux densities. Figure 3.13 shows the flux density ratio of 845 MALS and NVSS sources, and Figure 3.14 shows the flux density ratio of 1949 MALS and RACS sources. Only sources that are matched to a single source are used, and flux densities have been converted to the MALS rest frequency (1.27 GHz) assuming $\alpha = -0.75$. In both figures the corrections for the residual primary beam effect have properly re-scaled the flux densities, as the median flux density ratio (blue line) stays largely consistent with distance from the pointing centre, but a residual effect is left towards the edges of the image. In Figure 3.13 it stands out immediately that there is a systematic offset between the MALS and NVSS flux densities, as MALS flux densities are 18% higher on average. The overall flux density scale offset is $S_{MALS}/S_{NVSS} = 1.18 \pm 0.26$. In contrast to NVSS, Figure 3.14 shows that the flux densities between MALS and RACS agree extremely well up to $\rho \sim 0.5$ degrees. The overall flux density scale offset is $S_{MALS}/S_{RACS} = 1.06 \pm 0.39$, with the 6% overall offset originating mostly from the outer parts of the primary beam.

Though the result from NVSS might indicate any number of problems that could cause the offset, the additional data from RACS rules out most of these assumptions. A likely source of uncertainty would be the assumption of spectral index; however, this would impact the RACS results far more significantly, with its rest frequency of 887 MHz. From the RACS flux density offset we can assume that the 6% offset stems mostly from the residual primary beam effect, but this can only explain part




Figure 3.12: Astrometric offsets to NVSS for all ten pointings combined. The median offsets are given by the grey dashed lines, with the grey area indicating the uncertainty. The majority of sources lie within a FWHM of the average minor axis of the clean beam. The data are binned where five or more sources occupy the defined bin area; otherwise, individual sources are shown.

of the offset seen in NVSS. If this offset is persistent, it points to a systematic effect affecting either NVSS or both MALS and RACS. Due to the relatively low sensitivity of the surveys, only about 10% of MALS sources are matched to a counterpart, which makes the error bars on the flux density offset measurement rather large. As such, the measured offset is within the uncertainty, preventing us from making any definitive statement on the flux density offset. Combined with the measurement from Section 3.2.5, the flux density scale of MALS does not currently significantly deviate from the expected value, but the offset seen here indicates that more data are needed.

3.4.3 General assessment of the complete catalogue

In Section 3.3 the individual pointings have been evaluated with respect to completeness, purity, and flux density recovery. Here these properties are assessed on the entire catalogue in order to understand the impact of these characteristics on the final data product.

Completeness

To statistically determine completeness of the data, we re-factored the completeness to make it consistent between different pointings. In order to achieve this, instead of expressing completeness as a function of flux density in units of Jansky, we show flux density in units of σ_{20} as defined in Section 3.2.6. As we showed in Section 3.3.2, completeness for point sources scales linearly with local noise, and thus should be 0.2 at $5\sigma_{20}$ for all pointings. Though σ_{20} has units of Jy/beam, for




Figure 3.13: Ratio of flux densities of the sources in MALS compared to their NVSS counterparts as a function of distance from the pointing centre (ρ). The running median flux density ratio of the analytical primary beam correction both with (blue line) and without (red line) the corrections made in Section 3.4.1 are shown, as well as the running median flux density ratio of the holographic wideband primary beam correction (green line).

point sources integrated flux density and peak flux density are in principle equal, which means in this definition σ_{20} has the same value in Jy. Figure 3.15 shows that in terms of σ_{20} , pointings have very similar completeness curves, which allows us to combine the individual pointings and evaluate completeness for the whole survey, as indicated by the black combined completeness curve.

Combined completeness is also assessed as a function of separation from the pointing centre using only point sources, and as a function of major axis of the source using resolved sources. Both are shown in Figure 3.16. Combining the completeness from all the pointings gives enough statistical power to paint a clear picture of how the completeness is dependent on these variables. A clear relation is shown between completeness and distance from the pointing centre. The major difference between individual pointings seems to be the influence of the central source on the completeness. These differences are however extremely well modelled by the radially averaged rms noise (see Figure 3.8), indicating that completeness is related to the local noise. The right plot in Figure 3.16 shows that there is a power law decrease in completeness for larger sources. This was already suggested in Figure 3.9, but with the combined catalogues we have enough number counts to fully cover the space.

Flux recovery

In evaluating the individual pointings in Section 3.3.2, faint sources on average had higher measured flux densities compared to input flux densities. To fully assess this effect, we combine the flux recovery statistics from all pointings in Figure 3.17. In the combined statistics the effect is clearer, with bins farther away from the flux density ratio of unity being occupied with on average 10 sources




Figure 3.14: Ratio of flux densities of the sources in MALS compared to their RACS counterparts as a function of distance from the pointing centre (ρ). The running median flux density ratio of the analytical primary beam correction both with (blue line) and without (red line) the corrections made in Section 3.4.1 are shown, as well as the running median flux density ratio of the holographic wideband primary beam correction (green line).

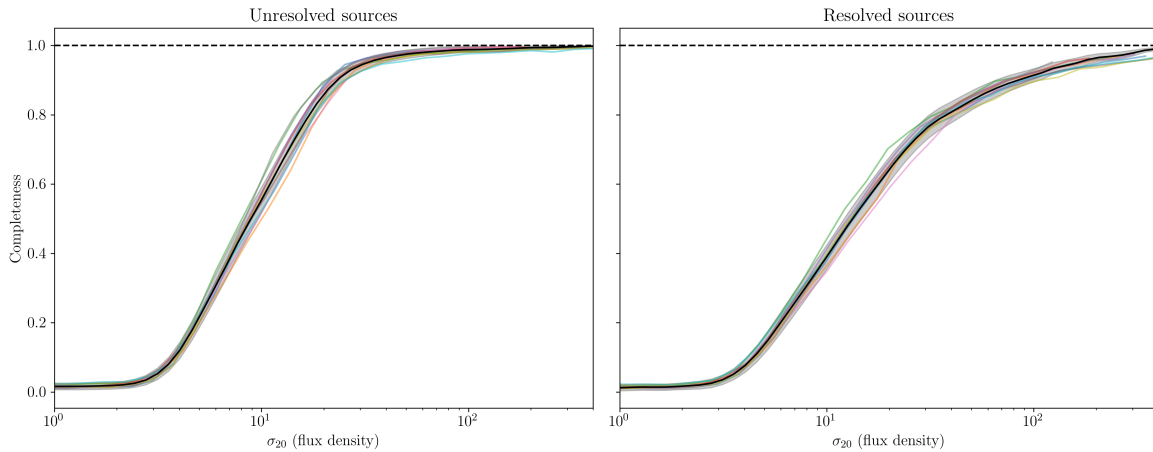


Figure 3.15: Completeness as a function of flux density for unresolved (left) and resolved (right) sources for the different fields, re-factored with σ_{20} and combined (black curves). Re-factoring the completeness curves to σ_{20} shows clearly that they are simply shifted with respect to each other, and we can define a unified completeness measure for the survey as a function of σ_{20} for both resolved and unresolved sources.




Figure 3.16: Combined source completeness as a function of distance from the pointing centre (unresolved sources, left) and major axis of the source (resolved sources, right). The left plot reflects the overall structure of the pointings, and shows that completeness is quite straightforwardly a radially averaged version of the noise structure as shown in Figure 3.3. Note that the flux density is normalised by σ_{20} . The right plot indicates a clear power law relation between the size of sources and the completeness, where larger sources are on average less complete.

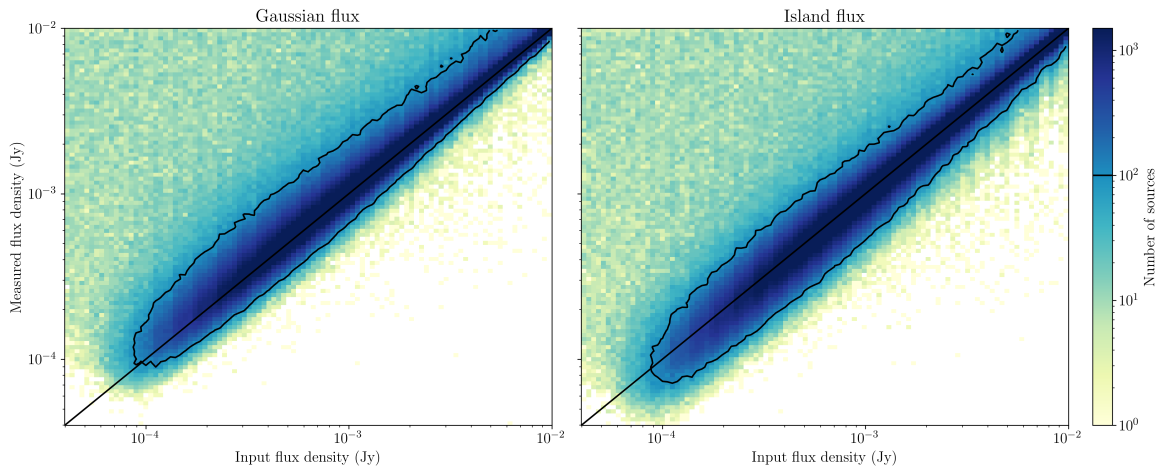


Figure 3.17: Input flux density plotted against the measured flux density for both Gaussian flux densities (left) and Island flux densities (right). The threshold of 100 sources per bin (black contour) shows quite clearly the bias present in Gaussian flux densities compared to island flux densities.




Figure 3.18: Purity of catalogues as a function flux density (left) and separation from the pointing centre (right). The fraction of false detections is indicated by the red line, both the total fraction (dotted line) and with removal of sidelobes (solid line). The open histograms show the number of sources detected in the pointings, with the filled histograms indicating the number of false detections. Though there seems to be no strong relation between flux density and purity, the number of false detections is strongly dependent on distance from the pointing centre, increasing both towards the centre and towards the edges of the pointing. Our criterion for identifying artefacts flags most of the false detections around the central source.

per bin. There is no visible dependence on flux density or distance from the flux density ratio of unity. Assuming a Poisson distribution of these bins with mean and variance $\lambda = 10$, we take all bins with fewer than 25 sources (5σ) to be part of this distribution. These bins combined contain 1.7% of all sources, indicating that this effect is rather small in terms of induced bias. Up to this point we have assumed that the flux density measured from the Gaussian fitting (the `Total_flux` column in the catalogues) best represents the flux density of the sources. In Figure 3.17 we compare the flux recovery between the Gaussian flux density and the integrated flux density from the island that the source occupies, where the contour indicates the threshold of 100 sources per bin. We see that across the board Gaussian flux densities are skewed towards higher values, where island flux densities remain symmetric around the input flux density. This is an effect that can significantly affect our catalogues, especially considering the increased number counts at lower flux densities. Consequently, in the analysis in the rest of the paper, we assume that the flux densities of sources are more accurately represented by the island flux density.

Purity

Combining false detections from all ten pointings, 241 sources are detected in the inverted images, making up 1.5% of the combined catalogue. As described in Section 3.3.2, our artefact identification criterion flags 44 of these, leaving 197 sources, or 1.2% of the combined catalogue. With the combined catalogue of false detections, we can investigate how purity is affected by other variables such as flux density and distance from the pointing centre. This will allow us to properly account for the purity of the catalogue, in order to not overestimate number counts. Given the variety of ‘source types’ seen in negative images and what counterparts we expect to see in the positive, the overall amount of false detections should be lower than the amount of sources seen in the negative image. In this sense, the purity is more appropriately an upper limit rather than a direct measure of false detections.




Figure 3.19: Ratio of total to peak flux as a function of signal to noise of both unresolved (blue) and resolved (red) sources in the combined catalogue.

Figure 3.18 shows the combined purity as a function of flux density and distance from the pointing centre. The left plot shows the purity as a function of flux density and shows that the fraction of false detections increases with higher flux density. This is largely a result of the overall number of sources decreasing at higher flux density, but does show that the flux density distribution of false detections is not the same as that of real sources. The lack of false detections at low flux densities shows that our 5σ detection threshold does not lead to a lot of spurious detections. It is noteworthy that more sources are flagged as artefacts at higher flux densities, indicating that artefacts around bright sources have higher flux densities on average. The right plot of Figure 3.18 shows a strong dependence of purity on the separation from the pointing centre, similar to the completeness. False detections increase near the central source because of strong artefacts, and there is a steady increase towards the edges of the primary beam. We see that our artefact selection criterion correctly picked out most of the artefacts originating from sidelobes of the central source, which dramatically increases the purity in the central portion of the image. Although the number of false detections restricts the statistical power of these results, the relations already show clear trends for the purity as a function of flux density and distance from the pointing centre that can be used when assessing number counts.

3.4.4 Resolved and unresolved sources

For the analysis of completeness in Sects. 3.3.2 and 3.4.3, we assumed that our catalogues are populated with both unresolved or resolved sources, and that these should be assessed separately. Figure 3.7 shows this distinction is warranted, as these sources types have very different completeness relations. If we want to apply this knowledge to real sources in the catalogue, we must have a reliable way of determining whether a source is resolved. We expect sources are resolved when their size exceeds

the size of the synthesised beam of the image; however, we must take the uncertainties introduced by noise in the image and fitting errors into account.

We determine source size by measuring the ratio between integrated flux density S and peak flux density S_{peak} of the source, which should be equal to one for an unresolved source. Figure 3.19 shows S/S_{peak} as a function of S/N, for both resolved and unresolved sources in our combined catalogue. Here and in subsequent usage of S/N we define it as the ratio between the peak flux density of the source and the local rms. Due to a combination of uncertainties, unresolved sources follow a log-normal distribution in S/S_{peak} (Franzen et al., 2015), and thus a normal distribution in $R = \ln(S/S_{peak})$ with mean 0 and standard deviation σ_R ,

$$\sigma_R = \sqrt{\left(\frac{\sigma_S}{S}\right)^2 + \left(\frac{\sigma_{S_{peak}}}{S_{peak}}\right)^2}. \quad (3.14)$$

We take both the uncertainties σ_S and $\sigma_{S_{peak}}$ as the sum in quadrature of their errors as determined by PyBDSF and a calibration error of 3%. The magnitude of the error in calibration is motivated by assessing Gaussian fits of bright unresolved sources and the flux density offset determined in Section 3.2.5. Using these quantities, the compactness criterion is then

$$\ln\left(\frac{S}{S_{peak}}\right) > 1.25\sigma_R. \quad (3.15)$$

The factor of $1.25\sigma_R$ encloses 95% of sources below $S/S_{peak} = 1$, so with the symmetry of the Gaussian distribution, 95% of all unresolved sources should be correctly identified with this criterion. As can be appreciated in Figure 3.19, according to this metric, 50% of all sources in the combined catalogue are resolved.

3.4.5 Catalogue columns

In the final catalogue, the majority of the columns are preserved from the PyBDSF source catalogues. Additional columns are however added in subsequent steps where required. Our aim is to only add information, and not remove any. This means that, for example, sources can be flagged as artefacts, but will still be present in the catalogue. When performing additional corrections on source flux densities and spectral indices, the correction factors are inserted into the catalogues so that the original values can be easily reproduced. The catalogue has 48 columns in total. Several lines of the catalogue are shown in Table 3.B.1 as an example.

Pointing_id - The ID of the pointing where the source has been found formatted as PT-JHHMM±HHMM.

Source_name - Name of the source, following IAU convention, formatted as JHHMMSS.S±HHMMSS.S with prefix MAL.S.

Source_id - Source ID as assigned by PyBDSF.

Isl_id - Island ID as assigned by PyBDSF.

RA and DEC (and errors) - The J2000 position of the source, defined as the centre of the composite Gaussian of the source, and associated errors.

Sep_PC - Distance of the source from the pointing centre.

Total_flux (and error) - Total flux density of the source and associated error.

- `Peak_flux` (and error) - Measured peak flux of the source and associated error.
- `Spectral_index` (and error) - Spectral index of the source, measured from the spectral index image, and associated error.
- `Spectral_index_correction` - The correction factor for residual primary beam effects on the spectral index of the source.
- `RA_max` and `DEC_max` (and errors) - Position of maximum intensity of the source and associated errors.
- `Maj`, `Min`, `PA` (and errors) - FWHM of the major axis, minor axis and position angle of the source fit by `PyBDSF` and associated errors.
- `DC_Maj`, `DC_Min`, `DC_PA` (and errors) - FWHM of deconvolved major axis, minor axis, and position angle, and associated errors.
- `Isl_Total_flux` (and error) - Total integrated flux of the island in which the source is located, and associated error.
- `Isl_rms` - Average background rms noise of the island in which the source is located.
- `Isl_mean` - Average background mean value of the island in which the source is located.
- `Resid_Isl_rms` - Average residual background rms noise of the island in which the source is located.
- `Resid_Isl_mean` - Average residual background mean value of the island in which the source is located.
- `S_Code` - Value generated by `PyBDSF` indicating whether a source is: fit by a single Gaussian ('S'), fit by multiple Gaussians ('M'), or one of multiple sources on the same island ('C').
- `N_Gaus` - Number of Gaussian components fit to the source.
- `Resolved` - Boolean indicating whether the source is resolved according to the metric defined in Section 3.4.4.
- `Flag_Artifact` - Boolean indicating whether the source is a likely artefact according to the criterion described in Section 3.3.1.
- `RA_mean` and `DEC_mean` - Mean intensity weighted position of all pixels of the island in which the source is located, measured if a source is fit with multiple Gaussians.
- `Cutout_Spectral_index` - Intensity weighted average spectral index of all pixels of the island in which the source is located, measured if a source is fit with multiple Gaussian components.
- `Cutout_Total_flux` - Total flux density of all pixels above the island threshold, measured if source fit with multiple Gaussian components.
- `Cutout_flag` - Flag assigned to cutout in certain conditions: the mean position falls outside the island ('M'), the position of the brightest pixel does not correspond to the maximum position measured by `PyBDSF` ('C'), the difference between `Cutout_total_flux` and `Isl_Total_flux` is more than 20% ('F').
- `Cutout_class` - Classification assigned at visual inspection as described in Section 3.3.1, indicating whether a source is well described by the Gaussian model ('G'), is better described by the island characteristics ('I'), is better described by a single Gaussian component ('P'), or an artefact ('A').

3.5 Source characteristics

When considering source counts in the radio regime, extra care must be taken in understanding the population of sources that is being probed. Depending on observing frequency and flux density,




Figure 3.20: Fractions of sources fit with varying numbers of Gaussians as a function of flux density. The number of components fit to sources increases steadily towards high flux density, but flattens out around 100 mJy, which can be caused by sources splitting up at these flux densities.

different source populations may appear in the sample. Our reference for source counts are the SKADS simulations, as the simulated sample is built up by different source populations. Given the theoretical noise limit of $10 \mu\text{Jy}/\text{beam}$, we can expect to detect sources down to $50 \mu\text{Jy}$. As shown in Wilman et al. (2008), the radio population is dominated by AGN above flux densities of 1 mJy, while below that star-forming galaxies start to make up a significant fraction of the source counts. There are several important distinctions between these source types that can influence source counts and a dipole measurement. Among them is source morphology, as multi-component sources can easily be mistaken for multiple separate sources, biasing number counts. Star-forming galaxies are primarily found at lower flux densities and can be morphologically described by a single component, such that we expect them to appear as faint point sources in our fields. Consequently, these sources can be easily counted as they are unambiguously unique sources and are thus statistically independent. At higher flux densities however, some sub-classes of AGN, such as Fanaroff-Riley type I (FRI, core-dominated) and type II (FR II, lobe-dominated) sources (Fanaroff and Riley, 1974), can boast extended structures that can complicate automated source finding methods.

3.5.1 Extended sources in PyBDSF

PyBDSF operates by fitting Gaussians to sources, which is an effective method for most radio sources, but breaks down in sources with more complex structure. PyBDSF offers multiple ways of improving the fit to extended sources, as specified in Section 3.3, but this chiefly improves detection of extended and more diffuse sources. To ensure that complex sources are accurately fit by a combination of Gaussians, we employ a special recipe for these types of sources, which make up 8% of the all sources

found in the fields. These sources are flagged in our workflow for visual inspection upon which the Gaussian fit is assessed as described in Section 3.3.1.

However, automated source finding algorithms will only recognise objects as a single source if they are closely connected, and will therefore fail on a subset of sources. This effect is strongest for FR II sources, as increased luminosity in the lobes makes them appear as separate radio sources. Source association is one of the outstanding problems in radio astronomy, and beyond the scope of this paper. Instead, we try to characterise this effect and how it might influence source counts. As the components of these sources are not statistically independent, they will naturally bias source counts. To get an estimate of how complex sources are fit by PyBDSF, we look at the number of components fit to sources as a function of flux density in Figure 3.20. We see a steady increase in the number of Gaussian components at higher flux densities; however, at $\gtrsim 100$ mJy the increase in components flattens out. This may indicate that at these flux densities we are seeing all the emission from these sources, and this is the ‘true’ distribution of components. However, an alternative explanation is that at these flux densities some extended sources no longer have connecting emission and are not recognised as single sources anymore. In this case the number of components keep increasing, but individual components will split off and be detected as separate sources, effectively keeping the number of components per source the same. Another indication of this happening might be seen in the differential source counts in Figure 3.24, where there is an excess in source counts above 100 mJy, relative to the expected values.

To get an alternative measure of this, we look at FR II galaxies in SKADS and the separation of their components. For an upper limit estimate on how many sources we expect to split up, we count components as separated when the distance between them exceeds $6.5''$, which is the minor axis of the average clean beam. With this, 11% of sources in the range 10 mJy - 1 Jy have two separated components. Furthermore, 6.3% and 24% of sources in the ranges 10-100 mJy and 100 mJy-10 Jy, respectively, have three separated components. Outside of these ranges the fractions are negligible. Doubling the distance required for separation mostly exchanges the amount of triple component sources for double component sources in the range 100 mJy - 1 Jy. The amount of triple component sources does not change in the range 1-10 Jy, indicating that the brightest sources are also the largest and most likely to separate. If we define the excess fraction of sources detected as $f_n = \frac{\tilde{n}-n}{n}$, where \tilde{n} is the amount of sources detected counting separate components and n the actual amount of sources, $f_n = 0.4$ at 10-100 mJy, $f_n = 0.9$ at 100 mJy-1 Jy, and $f_n = 0.7$ at 1-10 Jy. The values of f_n are given for each of the flux density bins used to determine number counts in Table 3.3. The maximum value f_n can take is 2, when a bin is entirely occupied by sources with three separate components. From this we can conclude that this effect is more important at higher flux densities, and is most significant at flux densities $\gtrsim 100$ mJy, which contains only a very tiny subset (102 sources, 0.6%) of the full catalogue.

3.5.2 Spectral indices

Due to different emission mechanisms and sources of emission, there can be differences in spectral index distribution between star-forming galaxies and AGN. Additionally, the Doppler shift observed as a consequence of the motion of the observer induces a change in observed flux density that depends on the spectral index of the source. Thus, the spectral indices of sources influences the magnitude of the radio dipole. In general, dipole studies assume a single value for spectral index based on the physics of synchrotron emission, $\alpha \approx -0.75$ near 1 GHz (e.g. Rubart and Schwarz, 2013; Tiwari and




Figure 3.21: Distribution of spectral indices of MALS sources. Left: MALS spectral indices (black) compared to AGN (red) and star-forming galaxies (blue) from SKADS as a function of flux density. Right: MALS spectral indices as a function of flux density, sources with S/N above 20 are coloured by S/N. The median value and error of spectral indices of different flux bins are indicated by red error bars, indicating that at lower flux densities spectral indices tend towards lower values.




Figure 3.22: Offsets of spectral indices measured from the wideband MTMFS images with respect to spectral indices derived from processing the full bandwidth using 15 individual SPW images from Deka et al. (2023) as a function of flux density. The spectral indices are calculated by using the SPW2 and SPW9 images. Binned median offsets are shown (blue), along with the offsets without the correction applied in Section 3.4.1 (red), showing that the spectral indices are properly corrected.

Jain, 2013; Siewert et al., 2021). Measuring the spectral index of sources generally requires either large bandwidth or measurements at different frequencies, which in turn requires high S/N to ensure that sources are detected at both ends of the frequency range. With the large bandwidth (802.5 MHz) of MALS we are able to create spectral index images, as described in Section 3.2.3, and measure spectral indices of sources in the catalogue.

To compare the source types, we look at the spectral indices of all sources with respect to flux density. While we cannot completely separate these source types based on flux density or spectral index, these source populations are labelled accordingly in the SKADS simulated sample (Wilman et al., 2008). Figure 3.21 shows the distribution of spectral indices measured compared to the populations in the SKADS simulated sample. Here, the AGN are separated into FRI (orange) and FRII (red) sub-populations, as are star-forming galaxies separated into ‘starburst’ (light blue) and ‘normal’ (dark blue) galaxies. It is noteworthy that the spectral index distribution of AGN boasts two peaks, corresponding to lobes at $\alpha \approx -0.75$, and cores at $\alpha \approx -0.25$. These peaks are respectively dominated by FRII and FRI galaxies, but there is cross-contamination present as we have plotted source components rather than combined source characteristics.

The right plot of Figure 3.21 highlights the MALS sources with high S/N, which are more likely to be detected across the full band and thus have more reliable spectral index measurements. Though less tight, the peak at $\alpha \approx -0.75$ is present in the MALS sample. The distribution is broad enough to have mixed with the peak $\alpha \approx -0.25$, so this second peak is possibly lost in the data. To see if we can retrieve these two populations, we take all sources with $S/N > 20$ and assert that the spectral index of the sources above and below $\alpha = -0.5$ represent FRI and FRII sources, respectively. With this definition, there are 626 FRI sources, of which 54% are resolved and 16% are fit with multiple Gaussian components. The brightest of these can be appreciated in Figure 3.A.1, showing mostly point sources or sources where core emission dominates. FRII sources are more numerous, with 2581 present in the catalogue, of which 75% resolved and 30% fit with multiple Gaussian components. A set of the brightest FRII are also shown in Figure 3.A.2, showing more sources with two or more components representing radio lobes. These results show that the expected dichotomy in morphology between these sources is indeed present, with FRII sources being more likely identified as extended and/or resolved.

At lower flux densities there appears to be a discrepancy between measured and theoretical spectral indices. We see in Figure 3.21 that not only is there a wider distribution in the MALS data, spectral indices are steeper at lower flux densities compared to the SKADS sources. The median spectral index for sources with $S > 1$ mJy is $\alpha = -0.76$, while for sources with $S < 1$ mJy it is $\alpha = -1.17$. Because of the high sensitivity that is needed, spectral indices are not commonly measured at lower flux densities. Looking at deep field surveys however, we see that this result is inconsistent with the spectral indices found in the XMM-LSS/VIDEO deep field (Heywood et al., 2020), where it is found that at lower flux densities spectra flatten out. The $S/N > 20$ sources shown in the right plot of Figure 3.21, though increasing in spread at lower flux densities, are not affected by the same bias.

We further investigate the bias seen in low S/N sources, and verify the corrections made in Section 3.4.1. We compare our spectral indices to those generated by comparing flux densities in SPW2 (1.0 GHz) and SPW9 (1.38 GHz) of the same MALS data by Deka et al. (2023). Smolčić et al. (2017) find a discrepancy between the spectral indices generated by MTMFS deconvolution and those generated by comparing flux densities at different frequencies, so we make the same comparison in Figure 3.22, showing the offset between the MTMFS and SPW derived spectral indices. The median offsets for both corrected (blue) and uncorrected (red) are shown, indicating that spectral indices have




Figure 3.23: Differential source counts from the SKADS simulations. The complete SKADS sample is shown (black), as well as the sample extracted from SKADS and injected into the images (blue). The uncorrected number counts (beige) indicate the sources detected by the source finding routine, which are then corrected with the rms noise coverage (green).

been properly corrected for residual primary beam effects. Though the offset trends negatively at lower flux densities, it is well within the uncertainties. Overall, the offset between spectral indices is -0.06 ± 0.92 (0.16 ± 0.94 without corrections) for all sources, and -0.01 ± 0.78 for $S/N > 20$ sources, agreeing well between the catalogues. There is no systematic effect seen of the corrected spectral indices with respect to the distance to the pointing centre, indicating no residual primary beam contribution. Deka et al. (2023) observe an overall flattening at low S/N compared to our overall steepening, creating a discrepancy that is clearly showing at lower flux densities. Overall spectral indices appear to be reliable down to mJy flux densities, or S/N of 20. Considering the flux densities of these sources, it is unlikely that many star-forming galaxies are included in the high S/N sample, precluding an analysis of these sources.

3.5.3 Number counts

Now it is left for us to assert that we have the necessary number counts for a dipole measurement. Extrapolating from the combined catalogue of the ten pointings, a catalogue of the first 391 pointings is expected to carry $\sim 650,000$ sources, enough to produce a dipole estimate if most sources can be used. However, because our pointings are inhomogeneous, both in terms of internal structure as well as with respect to other pointings, we have to assess to which extent this affects number counts and whether corrections can be made to homogenise the catalogues. The most common method of comparing number counts to other surveys or simulations is to compute differential source counts, which describes the number of sources dN within a given flux density bin $S + dS$ per steradian on the




Figure 3.24: Differential source counts from MALS, uncorrected (beige circles) and corrected with completeness (purple hexagons) taking into account whether a source is resolved or unresolved. Lastly, source counts are also corrected with rms noise coverage (green squares), which for the lowest flux bins goes to zero, causing solutions to diverge. This is all compared to the SKADS source counts, both for the central frequency of 1270 MHz (black), and for the full frequency range (grey area), and to the source counts derived from the MeerKAT DEEP2 image and NVSS from Matthews et al. (2021b) (white diamonds). The number counts are tabulated in Table 3.3.

sky. This is usually multiplied by $S^{5/2}$, which would yield a flat curve in a static Euclidean Universe (Condon and Ransom, 2016).

To verify that we would be able to retrieve the correct number counts, we repeated the experiment carried out in Section 3.3.2, injecting and retrieving sources in the residual images. To simulate a realistic physical distribution of sources, we cut out an area equal to the size of the pointing from the SKADS simulated catalogue. We repeated this experiment for every pointing, each time choosing a random position in the SKADS sample as pointing centre. Figure 3.23 shows the differential number counts for all the stages of the experiment. The reference sample from the SKADS simulations (black) represents the full 10×10 degree area simulated in Wilman et al. (2008). Out of the full sample, we cut out ten pointings with the same sky area as the MALS pointings that are injected into the residual images (blue). We performed our source finding routine on these images and saw what number counts we could retrieve. Figure 3.23 shows that below a few millijansky flux densities, detected source counts begin to fall off (pink), indicating the limit of 100% completeness for the full catalogue. These number counts are normalised by the area coverage of the pointings; however, we can make a simple correction based on the fact that the area coverage is not constant between different flux bins due to varying rms noise. The actual area covered in a certain flux bin $S + dS$ can be obtained by taking the rms noise coverage (as shown in Figure 3.4) assuming a detection limit of 5σ . This basic correction

(green squares) produces correct number counts down to 100-200 μJy , showing that we can account for completeness of the catalogue down to this flux density. Below this, we reach the absolute sensitivity limit of the pointings, as the rms noise coverage is so low that it produces diverging results. Above 100 mJy, results are more scattered, mainly because of low number counts at these flux densities and the smaller sky coverage of our pointings (35.7 sq. deg.) compared to SKADS (100 sq. deg.).

Having shown that we can reproduce number counts for a large range of flux densities, we measure differential number counts for the combined MALS catalogue. Both corrected and uncorrected differential source counts are tabulated in Table 3.3 and shown in Figure 3.24, where they are compared to the SKADS simulated sample. We apply a completeness correction (purple hexagons) using either unresolved or resolved completeness based on whether the source is classified as such using the criterion from Section 3.4.4. Once again we also correct number counts using the rms noise coverage (green squares). For comparison, number counts from the full SKADS simulated sample are also shown, both for the central frequency of 1.27 GHz and for the full frequency band, accounting for the fact that for most sources flux density is not equal across the band. The number counts derived from the MeerKAT DEEP2 image and NVSS by Matthews et al. (2021b) are also shown (white diamonds). Error bars are computed taking into account Poisson uncertainties as well as source clustering following Heywood et al. (2013). The same data from Figure 3.24 is tabulated in Table 3.3, showing the numerical values of flux bins, raw number counts, and uncorrected and corrected differential number counts. Though it is not taken into consideration in Figure 3.24, the number of false detection per bin is included in Table 3.3. In all cases the number of false detections is smaller than the uncertainty on the number counts, from which we infer that the purity is a small factor compared to completeness of the catalogue.

As for the simulated sample, the corrections appear to hold down to 100-200 μJy , after which we reach the sensitivity limit and solutions diverge. Until solutions diverge, the completeness and rms noise coverage corrections produce similar results, indicating that a simple sky coverage correction performs well given the ease with which it can be generated. Given the agreement between the completeness and rms noise coverage corrections, we can safely say that the major contributor to completeness is the inhomogeneous sky coverage of the pointings. This results in the dependence of completeness on the local noise seen for point sources in Figure 3.8. With these corrections applied, Figure 3.24 shows a wide range of flux densities where the differential number counts deviate from the expected values. At high flux densities (>100 mJy), we see an increase in sources that can for some part be attributed to the central sources in the images, as well as single sources being classified as multiple sources, as described in Section 3.5.1. However, we see that number counts, except for the range 20 mJy - 200 mJy, are higher across the board than what we might expect given theoretical predictions. There is some evidence that the SKADS simulations underestimated the number of star-forming galaxies, causing lower counts at low flux densities compared to what is seen in nature (e.g. Hale et al., 2023). Though this can explain an offset at the lowest flux densities, this effect would only be significant up to millijansky flux densities, whereas our number counts are higher up to an order of magnitude above that. An alternative explanation is that, with its selection of high flux density sources as pointing targets, MALS is probing overdensities, which naturally boosts the number of sources in the pointings. Lastly, such an effect can also be produced by a systematic overestimation of flux densities, which is an option that cannot be ruled out at this stage. We expect that this offset might also be caused by low number statistics, and may disappear once more data are added.

Table 3.3: Differential source counts of MALS, including corrected counts using rms coverage, completeness of unresolved sources and completeness of resolved sources. Counts are normalised for the sky area of 35.7 degrees. Raw counts (N) and number of false detections in (N_{false}) per bin are also given. The excess fraction of sources due to separated components f_n is also given, based on a separation distance of $6.5''$ of FR II sources in SKADS.

S	S_{mean}	N	N_{false}	Sky coverage	$S^{5/2} \frac{dN}{dS}$	Corrected $S^{5/2} \frac{dN}{dS}$	Corrected $S^{5/2} \frac{dN}{dS}$	f_n
(mJy)	(mJy)			(sq. deg.)	($\text{Jy}^{3/2} \text{sr}^{-1}$)	Rms coverage ($\text{Jy}^{3/2} \text{sr}^{-1}$)	Completeness ($\text{Jy}^{3/2} \text{sr}^{-1}$)	SKADS
0.1 - 0.13	0.11	850 ± 33	4	2.2	0.377 ± 0.015	6.06 ± 0.24	2.90 ± 0.11	0.0
0.13 - 0.17	0.15	1254 ± 43	11	5.3	0.816 ± 0.028	5.51 ± 0.19	3.93 ± 0.13	0.0
0.17 - 0.22	0.19	1575 ± 51	13	8.7	1.50 ± 0.05	6.16 ± 0.20	5.23 ± 0.15	0.0
0.22 - 0.28	0.25	1708 ± 54	12	12.5	2.39 ± 0.076	6.82 ± 0.22	6.47 ± 0.18	0.0
0.28 - 0.36	0.32	1558 ± 50	20	16.6	3.20 ± 0.10	6.89 ± 0.22	6.91 ± 0.19	0.0
0.36 - 0.46	0.41	1393 ± 46	22	20.9	4.21 ± 0.14	7.18 ± 0.24	7.58 ± 0.21	0.0
0.46 - 0.6	0.53	1257 ± 43	19	25.4	5.57 ± 0.19	7.83 ± 0.27	8.76 ± 0.25	0.0
0.6 - 0.77	0.69	1053 ± 38	15	23.3	6.85 ± 0.25	10.5 ± 0.4	9.67 ± 0.30	0.0
0.77 - 1	0.89	881 ± 34	16	31.6	8.41 ± 0.33	9.49 ± 0.37	11.1 ± 0.4	0.0
1 - 1.3	1.1	769 ± 31	10	30.9	10.8 ± 0.5	12.5 ± 0.5	13.4 ± 0.5	0.0
1.3 - 1.7	1.5	635 ± 28	7	34.4	13.1 ± 0.6	13.6 ± 0.6	15.5 ± 0.6	0.0
1.7 - 2.2	1.9	503 ± 24	4	35.0	15.2 ± 0.7	15.5 ± 0.8	17.3 ± 0.8	0.0
2.2 - 2.8	2.5	411 ± 21	7	35.3	18.2 ± 1.0	18.4 ± 1.0	20.3 ± 1.0	0.0
2.8 - 3.6	3.2	318 ± 18	2	35.5	20.7 ± 1.2	20.8 ± 1.2	22.5 ± 1.3	0.0
3.6 - 4.6	4.1	307 ± 18	3	35.6	29.3 ± 1.8	29.4 ± 1.8	31.3 ± 1.8	0.01
4.6 - 6	5.3	232 ± 15	6	35.7	32.5 ± 2.2	32.5 ± 2.2	34.3 ± 2.3	0.02
6 - 7.7	6.9	199 ± 14	1	35.7	40.9 ± 3.0	40.9 ± 3.0	42.7 ± 3.0	0.07
7.7 - 10	8.9	160 ± 13	1	35.7	48.3 ± 3.9	48.3 ± 4.0	49.6 ± 4.0	0.11
10 - 13	11	137 ± 12	1	35.7	60.7 ± 5.3	60.7 ± 5.3	60.7 ± 5.3	0.16
13 - 17	15	136 ± 11	0	35.7	88.5 ± 7.8	88.5 ± 7.8	88.5 ± 7.8	0.22
17 - 22	19	84 ± 9	1	35.7	80.2 ± 8.9	80.2 ± 8.9	80.2 ± 8.9	0.32
22 - 28	25	100 ± 10	0	35.7	140 ± 14	140 ± 14	140 ± 14	0.5
28 - 36	32	48 ± 6	0	35.7	98.7 ± 14.4	98.7 ± 14.4	98.7 ± 14.4	0.5
36 - 46	41	50 ± 7	0	35.7	151 ± 22	151 ± 23	151 ± 22	0.47
46 - 60	53	36 ± 6	0	35.7	160 ± 27	160 ± 27	160 ± 27	0.63
60 - 77	69	34 ± 5	0	35.7	221 ± 38	221 ± 38	221 ± 38	0.73
77 - 100	89	24 ± 4	0	35.7	229 ± 47	229 ± 47	229 ± 47	0.82
100 - 130	110	21 ± 4	0	35.7	294 ± 65	294 ± 65	294 ± 65	0.93
130 - 170	150	19 ± 4	0	35.7	391 ± 90	391 ± 90	391 ± 90	1.0
170 - 220	190	13 ± 3	0	35.7	392 ± 109	392 ± 109	392 ± 109	0.67
220 - 280	250	10 ± 3	0	35.7	443 ± 140	443 ± 140	443 ± 140	0.75
280 - 360	320	10 ± 3	0	35.7	650 ± 206	650 ± 206	650 ± 206	0.89
360 - 460	410	9 ± 3	0	35.7	859 ± 287	859 ± 287	859 ± 287	1.9
460 - 600	530	6 ± 2	0	35.7	841 ± 344	841 ± 344	841 ± 344	0.75
600 - 770	690	5 ± 2	0	35.7	1030 ± 460	1030 ± 460	1030 ± 460	0.5
770 - 1000	890	3 ± 1	0	35.7	906 ± 523	906 ± 523	906 ± 523	1.5
1000 - 1300	1100	4 ± 2	0	35.7	1770 ± 890	1770 ± 890	1770 ± 890	1.2
1300 - 1700	1500	1 ± 1	0	35.7	650 ± 650	650 ± 650	650 ± 650	0.75
1700 - 2200	1900	1 ± 1	0	35.7	955 ± 955	955 ± 955	955 ± 955	0.0




Figure 3.25: Demerit scores calculated for the 391 currently observed MALS pointings, using bright (> 100 mJy) sources retrieved from their corresponding surveys (NVSS and SUMSS). The dotted line indicates a quality threshold of $d < 15$ mJy based on the quality of the pointings inspected in this paper.

3.6 Towards the cosmic radio dipole

With a thorough assessment of the quality of the pointings described in this work we have the opportunity to extrapolate our findings to the larger survey of 391 pointings, both in estimating how many pointings will be needed for a dipole estimate, as well as how to effectively homogenise the catalogues. Given these results, there are however some questions and limitations that remain, and these will have to be addressed in later works. We first estimate the statistical power we can reach with the observed MALS pointings. As we determined in Section 3.2.6, the demerit score is a strong indicator for pointings with high noise and thus low number counts. Consequently, the demerit score allows us to make predictions about the quality of other pointings in the survey. Of the ten pointings we have investigated, we consider seven of them to be of good quality based on their noise and source count values. Since all MALS pointings are in the footprint of either NVSS or SUMSS, we use these surveys to match sources and calculate the demerit score to predict quality of the images. Based on this principle Figure 3.25 shows the distribution of the demerit scores of the first 391 observed MALS pointings. Defining a quality threshold of $d < 15$ mJy based on the ten pointings we have investigated here, we see that 322 of 391 pointings are below this threshold. The seven good quality pointings average ~ 2000 sources per pointing, meaning that 100 such pointings will result in approximately 2×10^5 sources. If we choose our pointings to properly cover the sky along the axis of the dipole, this is the minimum number of sources required for a 3σ measurement of the cosmic radio dipole, assuming an amplitude equal to that of the CMB dipole.

Based on the demerit scores of the ten pointings and those of the first 391 pointings, we also see that our ten pointings well represent the average pointing, and we can use the differential number counts to

extrapolate them to the rest of the survey. This allows us to take into account the completeness of the survey so far. As shown in Figs. 3.23 and 3.24, we get correct source counts down to 100-200 μJy . Naturally, this means that sources below that flux density cannot be included in a dipole estimate, which leaves 13,663 sources in the combined catalogue. However, the corrections are essentially compensating for the missing sources, yielding effectively $\sim 28,000$ sources down to 200 μJy , meaning even fewer than 100 pointings would again suffice for a 3σ dipole measurement.

In both cases we may expect that the first 391 observed pointings will yield around a million sources, which should yield a dipole measurement at a significance level of 6.5σ , assuming an adequate coverage of the dipole axis and a dipole amplitude equal to that of the CMB (Ellis and Baldwin, 1984). Though it is clear that MALS will deliver the number counts needed for a significant measurement of the cosmic radio dipole, the larger challenge is making a measurement while accounting for the systematics present in the survey. In order to thus successfully measure the dipole with MALS, the way forwards is to build on the corrections introduced in Section 3.5.3.

3.6.1 Compiling a homogeneous catalogue

We have effectively shown that we can make a unified description of the properties of the average MALS pointing, which should now allow us to homogenise the catalogue. Though the structure of the survey, with deep coverage over distinct patches of sky, appears to not lend itself especially well to large-scale cosmology, the fact that these pointings are all equal area by design allows for straightforward discretisation. In the simplest use case, each pointing of MALS can thus be treated as a single unit simply containing N number of sources. When measuring number counts over a full sky, inhomogeneities between the pointings induce higher order multipoles in the data that will spill over into a dipole measurement. To homogenise the data and get an unbiased estimate of the dipole, we must account for the individual differences between the pointings and calculate the corrected effective number counts N_{eff} .

To get to effective number counts, the starting point is the corrections to number counts as shown in Figure 3.24 that are seen to largely compensate for incomplete catalogues. We can extend this treatment by assigning a ‘completeness factor’ to each individual source. As discussed in Section 3.5.3, the largest contributor is the incomplete sky coverage, which is well modelled with the rms noise coverage of the pointings. Immediately we can disentangle the completeness into a detection probability, $P(det)$, and sky coverage, Ω ,

$$Comp(S, Q_A, \sigma) = P(det|S, Q_A)\Omega(\sigma). \quad (3.16)$$

Here the detection probability has a power law dependence on source size Q_A (linear in log-log space, right plot of Figure 3.16), and assuming Gaussian errors on the flux density, detection probability should follow a Gaussian cumulative distribution function or similar sigmoid function¹⁴ as a function of flux density. Sky coverage is exclusively determined by the local rms σ of the source, encoding the noise structure of the pointing. This completeness factor can largely correct for the inhomogeneities present in the catalogues, but we can make it even more robust by including information on other

¹⁴ Sigmoid is the collective name of functions following ‘S’-shaped curves, which are well suited to describe the detection probability of data near the detection boundary.

investigated quantities. Using information on purity, we can define a ‘purity factor’,

$$Purity(S, \rho, pointing) = Purity(S, \rho)Purity_{pointing}, \quad (3.17)$$

which indicates how likely a source is to be a true positive. This depends on distance from the pointing centre ρ , flux density S (Figure 3.11) and has a multiplicative factor that indicates a purity level that is different per pointing (Figure 3.18). Finally, individual sources have associated uncertainties that can be used to weigh each source accordingly. An obvious choice is a weight based flux density, as the uncertainty in flux density σ_S is dependent on flux density S (Figure 3.17), which can be combined with a potential flux density scale error ΔS . We are not limited to one uncertainty factor, and a second choice that is relevant for a dipole measurement is the uncertainty in position $\sigma_{\vec{\phi}}$, which in absence of any systematics (as we see in Figure 3.12) is simply equal to the measured uncertainty $\Delta\vec{\phi}$. Combining all these measures, we can assign a weight to sources based on the quantities laid out,

$$w_S = \sigma_S^{-1}(S, \Delta S), \quad (3.18)$$

$$w_{\vec{\phi}} = \sigma_{\vec{\phi}}^{-1}(\Delta\vec{\phi}), \quad (3.19)$$

$$w_{eff} = w_{\vec{\phi}}w_S \frac{Purity(S, \rho, pointing)}{Comp(S, Q_A, \sigma)}. \quad (3.20)$$

$$= w_{\vec{\phi}}w_S \frac{Purity(S, \rho)Purity_{pointing}}{P(det|S, Q_A)\Omega(\sigma)}. \quad (3.21)$$

The effective weight factor, w_{eff} , fulfils a dual purpose in estimating the number counts. The completeness and purity factors correct the number counts, while the weights from the flux density and position errors then serve as a quality measure for each source, allowing us to measure the effective number density of the individual pointings,

$$N_{eff} = \frac{\sum_i^n w_{eff,i}}{\sum_i^n w_{\vec{\phi},i}w_{S,i}}. \quad (3.22)$$

3.6.2 Limitations and future prospects

With this prescription, the systematic effects that we have characterised can be accounted for when computing the number counts and estimating the dipole. There remain however some effects that have not been explicitly characterised that could influence a dipole measurement. By checking the corrections to number counts as we did in Section 3.5.3 on simulated data, we essentially calibrated the corrections on the SKADS sample, which as a simulation might not perfectly represent the number counts found in nature (this can be plainly seen in Figure 3.24, where SKADS number counts do not always agree with the counts from Matthews et al. (2021b)). While this can introduce an unknown error into the process, the error is expected to be in overall number counts, and therefore not directionally dependent. Similarly, an important aspect of MALS is the selection of the pointings, as every pointing has a bright radio source at the centre. Although pointings are distributed isotropically, this is not equivalent to a random selection as bright central sources are more likely to be embedded in overdensities. This effect seems to be very pronounced in our measured source counts already, which are larger than expected. Again, this effect is expected to be directionally independent, but whether

this is truly the case remains to be determined. Finally, the depth of MALS might be to its detriment when measuring a dipole, as the reached depth of $200 \mu\text{Jy}$ probes into the population of starburst and normal galaxies. The brightest sub-population of these fainter sources occurs at lowest redshifts, which exhibits stronger clustering than AGN. To what extent this affects a dipole measurement is explored in Bengaly et al. (2019), who perform several redshift cuts and a significant improvement is made with $z_{cut} = 0.1$ compared to including all sources. This is a more stringent cut than Blake and Wall (2002), who claim to eliminate local clustering effects with $z_{cut} = 0.03$. Although no direct redshift information is available from the MALS data, we will be able to investigate the effect of clustering due to nearby star-forming galaxies using photometric redshifts from current surveys such as the Dark Energy Camera Legacy Survey (DeCALs; Blum et al., 2016), and the Rubin Observatory Legacy Survey of Space and Time (LSST; LSST Science Collaboration et al., 2009) in the near future.

It is worth noting that an increase in number counts as seen in Figure 3.24 can also be (partially) caused by a systematic flux density offset. The results seen in Figure 3.13 hint to a systematic flux density offset with respect to NVSS, which is worth investigating with the larger MALS catalogue. If this offset turns out to be indeed significant, there are many possible explanations given that such a systematic effect could potentially be introduced at many points in the data processing pipeline. We have already shown in Section 3.4.3 that the island flux density from PyBDSF properly recovers the flux density from simulated sources, which verifies that the source finding step is not inducing a systematic flux density offset. To further narrow down the options, we cross-check our results with those from Deka et al. (2023), which use the same data, calibration pipeline and source finding strategy but show an overall agreement with NVSS. The most notable difference between these catalogues is the fact that we utilise the full band while Deka et al. (2023) only use individual SPWs with 50 MHz bandwidth. Such a difference between results from the full bandwidth and individual SPWs could point to systematic effects introduced in the imaging stage, as we model the emission with two Taylor terms in frequency as opposed to a single Taylor term in case of individual SPW images. Though the ten pointings explored here have already provided a wealth of statistics and insight into systematic effects, these will be further investigated with the full suit of MALS pointings.

With regards to a dipole measurement, there are still some questions that remain to be answered. Throughout this work we have used the estimate by Ellis and Baldwin (1984) of 2×10^5 sources properly distributed along the axis of the dipole. Though the MALS pointings properly cover the axis of the CMB dipole, if the direction of the radio dipole deviates from this, for example towards the northern hemisphere, the coverage of MALS pointings might not be adequate. Furthermore, the exact amount of sources needed for a dipole estimate can vary depending on this coverage and has been differently estimated in different works. Crawford (2009) claims 2×10^6 sources are necessary for a 3σ dipole estimate, which is an order of magnitude more than the number from Ellis and Baldwin (1984). Dipole studies using NVSS have generally reached 3σ significance with 3×10^5 sources (e.g. Singal, 2011; Rubart and Schwarz, 2013; Secrest et al., 2022). This is in closer agreement to the Ellis and Baldwin (1984) numbers; however, this significance is only reached because of the anomalously high dipole amplitude. Therefore, the significance with which the dipole can be measured ultimately depends on many factors, including sky coverage, number counts, dipole amplitude, frequency, and the employed estimator. Consequently, another important step towards measuring the cosmic radio dipole with MALS is defining an appropriate dipole estimator (see e.g. Siewert et al., 2021), which beyond the scope of this paper but will be explored in a future work.

3.7 Summary and conclusion

In this work we have presented a thorough analysis of the first ten deep continuum pointings of MALS (Gupta et al., 2016) and have compiled a catalogue with 16,307 sources covering 35.7 square degrees of deep radio sky. We set out to extensively analyse the properties of the first ten pointings of MALS, with the ultimate goal of measuring the cosmic radio dipole. To achieve a measurement of number counts unbiased by the inhomogeneities present between the MALS pointings, we characterised systematic effects that can influence such a measurement. This assessment of systematic effects in the ten pointings as presented in this work shows that these effects are for the most part predictable and can be properly accounted for. This will eventually not only benefit a dipole measurement, but all continuum science carried out with MALS. In the current literature on the cosmic dipole, there are many examples of systematic effects that limit the sensitivity of these estimates and could not be pushed further due to a lack of information on the inner workings of the surveys that were used. For MALS, we have a complete assessment of the inner workings of the survey, with insight and access into the processing pipeline. Looking forwards, we determine that 100 MALS pointings suffice for a dipole measurement. Paired with the analysis on source characteristics and counts in the pointings, we are poised to perform the most complete dipole estimate of a radio survey thus far.

Calibration and imaging of all the data was carried out through ARTIP. After imaging, we separately created spectral index images and performed primary beam correction averaged over the frequency range on these and the continuum images with a primary beam model derived from holographic measurements. We made an initial assessment of the calibration by checking the flux density scale of the calibrators and central sources, and find that flux densities are consistent with those reported in the literature. We investigated the quality of the images by looking at the rms noise maps created by PyBDSF. Measuring rms noise coverage shows that all pointings have similar noise structure, but overall noise levels are offset between the pointings. We quantified this offset with σ_{20} , which gives the noise level at 20% rms noise coverage for a pointing. To try to explain the difference in pointing quality, we calculated demerit scores for each pointing to estimate the contributions of bright sources to the noise. Though there is a correlation between the noise in a pointing and the demerit score, other factors play a part in introducing a scatter in this relation, especially for pointings with lower demerit scores. As σ_{20} directly describes the overall noise in the image, it is the quality measure of choice for the pointings. We performed source extraction on all the images using PyBDSF and converted PyBDSF catalogues to full Stokes I catalogues, extending them to include spectral indices and flagging artefacts. We considered sources fit with multiple Gaussian components to be potentially complex and visually inspected them to investigate how well the Gaussian fit describes these sources. We further assessed the quality of the individual pointings and how it affects source finding by measuring completeness, flux recovery, and purity. For completeness and flux recovery, we need to know the intrinsic properties of the sources in the images, so we created mock catalogues of sources from the SKADS simulated sample and injected these into the residual images of the pointings. We assessed completeness for unresolved and resolved sources separately. Using unresolved sources, we assessed completeness as a function of distance from the pointing centre, and for resolved sources we investigated completeness as a function of source size.

Combining the catalogues of the individual pointings, we corrected for residual primary beam effects in the flux densities and spectral indices sources originating from the frequency dependence of the primary beam. To check these corrections and other potential systematic effects, we cross-matched the catalogues with NVSS to check if positions and flux densities were consistent. There is no appreciable

astrometric offset; there is an 18% offset in flux density compared to NVSS, but this is still within the uncertainties. Using σ_{20} as a normalisation factor, we combined completeness measures from the individual pointings and find unified completeness relations that hold for all pointings. Combining flux recovery statistics from all pointings, we find that a systematic bias is present in the integrated flux densities from the fitted Gaussians in the catalogue. This bias is not present in the integrated flux densities of the islands that the sources occupy, making this quantity the logical choice of flux density for the analysis presented here. Combining purity from all pointings, we find that we can account for 20% of false detections with a suitable artefact identification scheme. The remaining false detections make up an increasing fraction of sources farther out from the pointing centre.

As the full catalogues are expected to be populated by various sources types, we assessed how this can influence source counts. Looking at the number of Gaussian components fit to sources, we see that the number of components needed to describe sources increases as a function of flux density until it stagnates at around 100 mJy. This implies that, around this flux, density sources separate and can be counted multiple times; however, as only a tiny percentage of sources are present at these flux densities, it is unlikely to bias the source counts. To further differentiate between source populations, we looked at the spectral indices of sources and find that we can reasonably separate core-dominated FRI from lobe-dominated FR II sources by making a cut at $\alpha = -0.5$. At low flux densities we find that spectral indices are much steeper than expected, which is likely caused by low S/N and inadequate modelling by the MTMFS deconvolution scheme.

Finally, to show that we can account for the systematic effects present in the catalogues, we calculated and corrected differential number counts for a set of simulated catalogues using the SKADS simulated sample. We find that a simple correction using only the rms noise coverage produces correct number counts down to 100-200 μ Jy. We then computed differential number counts for the full catalogue and corrected number counts using the combined completeness measures found for unresolved and resolved sources, as well as the rms noise coverage. Once again, corrections seem to hold down to 100-200 μ Jy. Comparing these number counts to the expected number counts from the SKADS sample, we see that our number counts are higher in the full range of probed flux densities. This is independent of corrections, so a likely explanation is that MALS is probing overdense regions, as it targets bright sources. The same effect can, however, be (partially) produced by a systematic flux density offset, which should be taken into consideration given the results on the flux density scale.

Using both the demerit score to predict the quality of other MALS pointings and the corrected number counts of this sample, we show that we will require 100 MALS pointings to reach the necessary number counts for a 3σ measurement of the dipole. Going further, we assert that we can assess the dipole on the level of individual sources using the information on flux density scale, completeness, purity, flux errors, and position errors. The precise implementation is left to later works, along with an exploration of viable dipole estimators.

3.A Cutouts of bright FRI and FRII sources

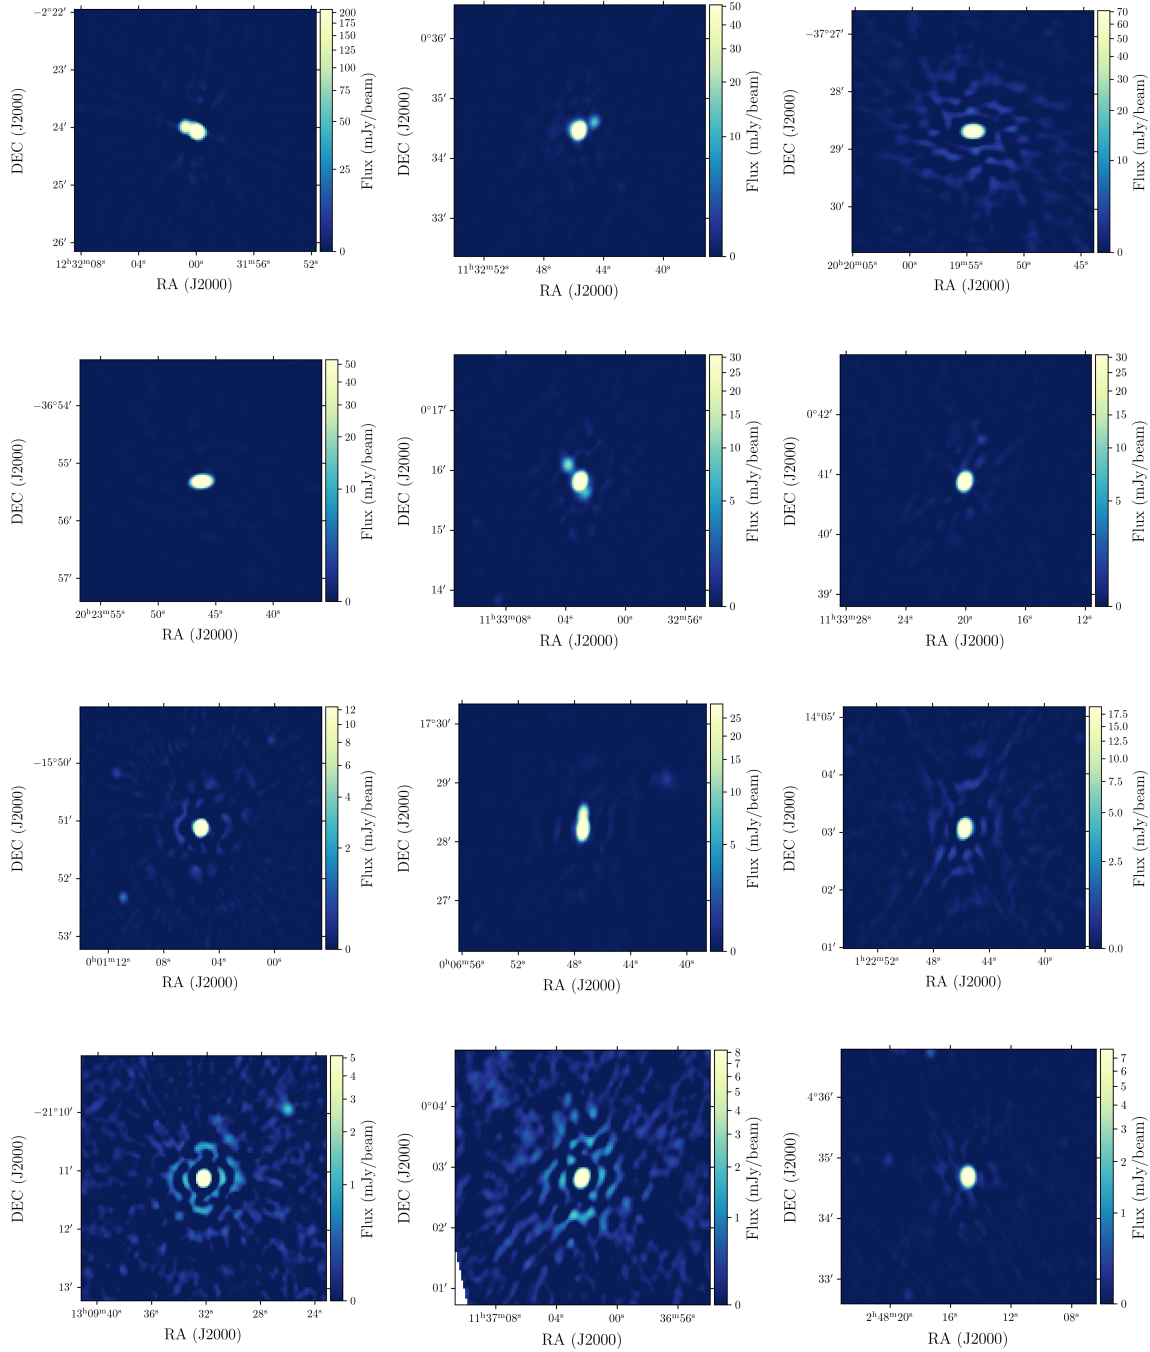


Figure 3.A.1: Selection of brightest sources in the combined catalogue with $\alpha \geq -0.5$. As discussed in Section 3.5.2, this range of spectral indices is expected to come from emission originating in AGN cores, which are generally point sources. The majority of this sample indeed is unresolved or can be seen to dominate the emission from their associated lobes.




Figure 3.A.2: Selection of brightest sources from the combined catalogue with $\alpha < -0.5$. As discussed in Section 3.5.2, this range of spectral indices is generally associated with synchrotron, which is the dominant emission mechanism in radio lobes. Though a number of sources in this sample are unresolved, many show the two component structure characteristic of FRIs or even more complex extended structure.

3.B Table of sources

Table 3.B.1: Example of ten rows of the final source catalogue, showing the catalogue structure. The full catalogue width spans all subtables below. The 48 columns are described in Section 3.4.5.

	Pointing_id	Source_name	Source_id	Isl_id	RA J2000, °	E_RA ° ($\times 10^{-4}$)
1	PT-J0001-1540	J000012.63-154312.0	1607	1618	0.052609	0.68
2	PT-J0006+1728	J000409.33+163842.5	1285	1331	1.038859	0.83
3	PT-J0126+1420	J012317.38+133309.1	1504	1530	20.822411	1.70
4	PT-J0240+0957	J023823.97+094637.3	888	904	39.599861	0.31
5	PT-J0249+0440	J024648.18+035955.5	1469	1483	41.700764	0.85
6	PT-J0249-0759	J024633.91-073732.5	2518	2529	41.641290	1.64
7	PT-J1133+0015	J113059.76+000708.0	708	717	172.749000	2.06
8	PT-J1232-0224	J122945.15-021418.1	507	517	187.438138	1.29
9	PT-J1312-2026	J130900.75-204924.6	2326	2334	197.253139	0.40
10	PT-J2023-3655	J202008.96-371033.3	2055	2067	305.037343	2.38

	DEC J2000, °	E_DEC ° ($\times 10^{-4}$)	Sep_PC °	Total_flux mJy	E_Total_flux mJy	Peak_flux mJy/beam	E_Peak_flux mJy/beam
1	-15.720001	0.94	0.36	0.409	0.050	0.273	0.022
2	16.645139	1.36	1.04	4.564	0.458	2.101	0.151
3	13.552529	1.65	1.06	19.007	1.561	2.288	0.121
4	9.777035	0.59	0.54	4.692	0.368	2.575	0.062
5	3.998738	1.61	0.98	1.418	0.231	0.985	0.102
6	-7.625700	2.18	0.83	0.421	0.113	0.285	0.049
7	0.118892	2.53	0.53	0.423	0.114	0.254	0.046
8	-2.238372	1.77	0.59	0.530	0.121	0.394	0.057
9	-20.823501	0.50	0.82	1.761	0.117	1.136	0.050
10	-37.175915	1.73	0.77	0.511	0.116	0.306	0.047

	Spectral_index	Spectral_index_correction	E_Spectral_index	RA_max J2000, °	E_RA_max ° ($\times 10^{-4}$)
1	-1.90		0.10	0.052609	0.68
2	-0.10		0.80	1.038859	0.83
3	-0.53		0.83	20.824116	1.70
4	-0.74		0.21	39.601394	0.30
5	0.46		0.72	41.700764	0.85
6	-0.94		0.51	41.641290	1.64
7	-1.51		0.21	172.749000	2.06
8	-2.36		0.25	187.438138	1.29
9	0.09		0.50	197.253139	0.40
10	-1.56		0.43	305.037343	2.38

Table 3.B.1: continued

	DEC_max J2000, °	E_DEC_max ° ($\times 10^{-4}$)	Maj "	E_Maj "	Min "	E_Min "	PA °	E_PA °
1	-15.720001	0.94	9.35	0.80	7.78	0.57	6.7	19.6
2	16.645139	1.36	14.55	1.16	10.66	0.70	171.8	11.2
3	13.551953	1.65	19.38	1.50	17.70	1.33	38.2	36.8
4	9.776665	0.59	14.21	0.51	9.25	0.25	80.3	3.8
5	3.998738	1.61	11.36	1.38	7.74	0.69	10.1	13.4
6	-7.625700	2.18	10.19	1.97	7.82	1.22	154.0	30.8
7	0.118892	2.53	11.83	2.42	8.42	1.34	146.1	26.1
8	-2.238372	1.77	9.72	1.53	7.90	1.04	162.6	31.5
9	-20.823501	0.50	9.35	0.43	8.03	0.33	17.9	13.0
10	-37.175915	1.73	12.52	2.08	9.91	1.38	72.2	30.8
	DC_Maj "	E_DC_Maj "	DC_Min "	E_DC_Min "	DC_PA °	E_DC_PA °		
1	5.60	0.80	4.16	0.57	31.1	19.6		
2	9.03	1.16	8.62	0.70	161.0	11.2		
3	17.71	1.50	15.05	1.33	64.8	36.8		
4	0.00	0.51	0.00	0.25	0.00	3.8		
5	6.93	1.38	3.53	0.69	22.8	13.4		
6	6.49	1.97	3.62	1.22	141.4	30.8		
7	8.15	2.42	4.45	1.34	133.2	26.1		
8	4.97	1.53	3.85	1.04	135.7	31.5		
9	6.02	0.43	4.12	0.33	44.7	13.00		
10	8.28	2.08	4.65	1.38	29.8	30.8		
	Isl_Total_flux mJy	EIsl_Total_flux mJy	Isl_rms mJy/beam	Isl_mean mJy/beam	ResidIsl_rms mJy/beam			
1	0.369	0.034	0.020	-0.004	0.007			
2	4.343	0.276	0.141	-0.021	0.121			
3	11.940	0.416	0.121	-0.010	0.121			
4	3.954	0.160	0.062	0.005	0.021			
5	1.231	0.136	0.098	-0.020	0.019			
6	0.349	0.060	0.047	-0.000	0.009			
7	0.384	0.064	0.043	0.003	0.003			
8	0.426	0.065	0.054	-0.001	0.006			
9	1.613	0.082	0.046	-0.019	0.035			
10	0.373	0.056	0.043	-0.008	0.007			

Table 3.B.1: continued

	ResidIsl_mean mJy/beam	S_Code	N_Gaus	Resolved	Flag_artifact	RA_mean J2000, °	DEC_mean J2000, °
1	-0.004	S	1	True	False	–	–
2	-0.002	S	1	True	False	–	–
3	-0.010	M	2	True	False	20.822423	13.552525
4	0.004	M	2	True	False	39.599830	9.777032
5	-0.019	S	1	True	False	–	–
6	-0.000	S	1	False	False	–	–
7	0.003	S	1	True	False	–	–
8	-0.001	S	1	False	False	–	–
9	-0.014	S	1	True	False	–	–
10	-0.008	S	1	False	False	–	–

	Cutout_Total_flux mJy	Cutout_Spectral_index	Cutout_flag	Cutout_class
1	–	–	–	–
2	–	–	–	–
3	10.482	-0.24	C	I
4	3.583	-0.75	C	P
5	–	–	–	–
6	–	–	–	–
7	–	–	–	–
8	–	–	–	–
9	–	–	–	–
10	–	–	–	–

Rise of the Bayesian estimators

The cosmic radio dipole: Bayesian estimators on new and old radio surveys

J. D. Wagenveld, H-R. Klöckner, and D. J. Schwarz

Astronomy & Astrophysics, 675, A72 (2023)

Overview

In this chapter, we present a novel set of Bayesian estimators to determine the cosmic radio dipole in preparation for a MALS measurement. Where in the previous chapter we focused on characterising and accounting for systematic effects with inside knowledge of the survey, we focus here on accounting for the systematic effects using only information that is already present in the catalogues.

For the estimation of the number count dipole, the sky is commonly divided into equal size cells, and the amount of sources is counted per cell. The expectation for an isotropic distribution of sources is that the counts-in-cell distribution of pixel values then follows a Poisson distribution. We introduced several estimators based on the principle of maximum likelihood estimation. For our most basic estimator, we assumed a dipole effect which affects source density as a function of location on the sky. We defined a second estimator, where we used the assumption that local variations in source density were directly related to the local rms noise. We modelled this with a power law relation between the local noise and the source density, and integrated this into the Poisson estimator. Finally, we defined a third estimator that could do a joint estimate on several catalogues at once, assuming that the catalogues had different source densities but the same dipole. For all of these estimators, the monopole(s) and the dipole direction and amplitude were free parameters. For the Poisson-rms estimator using local rms noise, the index x of the power law describing the relation between the noise and the source density was a free parameter as well.

To test these estimators, we used them to estimate the dipole on simulated data, as well as the Rapid Australian Square Kilometre Array Pathfinder (ASKAP) Continuum Survey (RACS, McConnell et al., 2020) and National Radio Astronomy Observatory (NRAO) Very Large Array (VLA) Sky Survey (NVSS, Condon et al., 1998) radio catalogues. RACS and NVSS cover the southern and northern sky respectively, and thus complement each other well in terms of sky coverage. Though the dipole has been measured many times with the NVSS catalogue, the RACS catalogue is less well explored in terms of dipole measurements. RACS however has measurements of the local noise in the catalogue,

making it an ideal candidate for a dipole measurement with Poisson-rms estimator.

We compared our results from the basic Poisson estimator to the more commonly used quadratic estimator, and found them to be consistent. Furthermore, measurements between NVSS and RACS were also found to be consistent in terms of dipole amplitude, both being around three times higher than the CMB dipole amplitude. In terms of dipole direction, RACS and NVSS were slightly offset from each other and the CMB dipole, they were however all consistent within 3σ . The result for the RACS dipole with the Poisson-rms estimator was deviating both in terms of amplitude and direction, indicating that some additional systematic effects were present that were not captured by the relation between the local noise and the source density. A combined parameter estimation using both NVSS and RACS resulted in a dipole estimate that perfectly aligns with the CMB dipole in terms of direction but with an amplitude that is three times as large. This result has a significance of 4.8σ , which is an unprecedented level of precision for radio sources, only matched by recent results using infrared AGN.

In a final experiment, we used the estimator for multiple catalogues on the NVSS, separating the parts of the catalogue where different VLA configurations were used into separate catalogues. We showed that the north-south anisotropy that dominates any dipole estimate for low flux density cuts can be mostly alleviated with this technique, which is a promising prospect for upcoming surveys that use a similar scheme such as VLASS.

This paper was originally published as Wagenveld et al. (2023b) in *Astronomy & Astrophysics* in 2023. It can be found in its original form in Appendix C. It is reproduced below with some minor adjustments to support the flow of this thesis. As the lead author of the paper, I wrote the manuscript and produced all the figures. I wrote the Bayesian analysis code that was released along with the paper, and performed the checks and analysis on the data.

4.1 Introduction

In Chapter 3 we presented a deep analysis of ten pointings from MALS, with a focus on mitigating biases that could affect a measurement of the cosmic radio dipole. Being able to account for systematic effects allows less strict flux density cuts to be made, increasing the homogeneity of the catalogues and the number of sources that can be used for a dipole estimate. This approach was made possible by having direct access to meta data and data products within the processing steps of the survey from calibration to imaging and source finding, which is commonly not the case in dipole studies. In the present work, we approach this problem from the outside, and show how dipole estimates can be improved with the information present in modern radio catalogues. We present new dipole estimators, constructing likelihoods for estimating dipole parameters that can take this information into account, as well as an estimator that combines different catalogues for an improved dipole estimate. Given the proper information, these estimators are able to account for systematic effects on number counts in radio surveys and remove the need to cut large amounts of data.

This paper is organised as follows. In Section 4.2, we describe the statistics of radio source counts and the dipole effect. In Section 4.3, we introduce the estimators that will be used to infer the radio dipole parameters using the data sets described in Section 4.4. The results obtained are given in Section 4.5. The implications of our results are discussed in Section 4.6, along with some caveats. In Section 4.7, we summarise the findings of our study.

4.2 Radio source counts and the dipole

The majority of bright sources at radio wavelengths outside of the Galactic plane are active galactic nuclei (AGN) and have a redshift distribution that peaks at $z \sim 0.8$ (e.g. Condon and Ransom, 2016). As such, radio sources are expected to trace the background, and should therefore comply with isotropy and homogeneity on the largest scales. Following the cosmological principle, the surface density of sources should therefore be independent of location on the sky. The naive expectation is a distribution of radio sources that are independent, identical, and point-like, which defines a Poisson point process. By discretising the sky into regions of finite size, the number of sources per region will follow a Poisson distribution. The probability density distribution

$$p(n) = \frac{\lambda^n e^{-\lambda}}{n!} \quad (4.1)$$

is entirely parametrised by the variable λ , which describes both the mean and variance of the distribution. In actual radio data, some deviations from a perfect Poisson distribution are expected due to clustering and the presence of sources with multiple components. The severity of these effects depends largely on factors such as survey depth, angular resolution, and observing frequency, and is therefore difficult to assess without a thorough analysis of the survey. Such an analysis is beyond the scope of this work, but would follow a structure similar to that of Siewert et al. (2020), who demonstrate for the LOFAR Two-Metre Sky Survey first data release (LoTSS DR1, Shimwell et al., 2019) that the distribution of source counts converges to a Poisson distribution if applying stricter flux density cuts. For the analysis presented in this work, we assume that the effects of clustering and multi-component sources are negligible on a dipole estimate.

The spectral features and number count relations of typical radio sources make them uniquely suitable for a dipole measurement. For most sources, the dominant emission mechanism at radio wavelengths is synchrotron radiation, the spectral behaviour of which is well described by a power law,

$$S \propto \nu^{-\alpha}, \quad (4.2)$$

with a characteristic spectral index α . For synchrotron emission, the typical value of α is around 0.75 (e.g. Condon, 1992), and this value has been assumed for most dipole studies at radio wavelengths (Ellis and Baldwin, 1984; Rubart and Schwarz, 2013; Siewert et al., 2021). Furthermore, the number density of radio sources follows a power-law relation with respect to the flux density S above which the counts are taken,

$$\frac{dN}{d\Omega}(> S) \propto S^{-x}. \quad (4.3)$$

The value of x can differ between surveys depending on the choice of flux density cut and frequency of the catalogue, but usually takes values of 0.75-1.0.

In the frame of the moving observer, given a velocity $\beta = v/c$, a systemic Doppler effect shifts the spectra of sources, which affects the flux density of these sources. Additionally, sources are Doppler boosted from the point of view of the moving observer, which further affects the observed flux density of sources. Depending on their angular distance from the direction of motion θ , the flux densities are shifted by

$$S_{obs} = (1 + \beta \cos \theta)^{1+\alpha} S_{rest}. \quad (4.4)$$

Thus, given a flux-limited survey of radio sources, more sources appear above the minimum observable flux in the direction of the motion, and less will appear in the opposite direction. Finally, relativistic aberration caused by the motion of the observer shifts the positions of sources towards the direction of motion, causing a further increase in number counts in the direction of motion,

$$\tan \theta_{obs} = \frac{\sin \theta_{rest}}{\beta - \cos \theta_{rest}}. \quad (4.5)$$

As the fluxes and positions of the sources are shifted, we observe the dipole as an asymmetry in the number counts of radio sources. Combining these effects to first order in β , and therefore assuming that $v \ll c$, shows the expected dipole amplitude for a given survey:

$$\vec{d} = \mathcal{D} \cos \vec{\theta}, \quad (4.6)$$

$$\mathcal{D} = [2 + x(1 + \alpha)]\beta. \quad (4.7)$$

As such, we can directly infer the velocity of the observer by measuring the dipole effect on the number counts of sources. However, this necessitates the assumption that the dipole is entirely caused by the motion of the observer. Given the observed discrepancy between the CMB dipole and the radio dipole, it would be equally appropriate to assume that part of the observed radio dipole is caused by a different (and as-of-yet unknown) effect.

Given a dipole characterised by Equation 4.7, different dipole amplitudes \mathcal{D} are expected to be seen depending on the data set being used. Though aberration always has an equal effect, the effect of Doppler shift is determined by the spectral index α of the sources, and both this and the Doppler-boosting effect depend on the flux distribution of sources, which is characterised by the power-law index x . Single values are most often assumed for these quantities, from which the expectation of the dipole amplitude can be derived. These quantities will differ at least between different surveys, and so we derive them for each survey separately. Although the entire flux distribution of any survey cannot be characterised with a single power-law relation, the dipole in number counts is caused by sources near the flux density threshold, making the power-law fit near this threshold the most appropriate choice for deriving x . For the entire range of frequencies considered in this work, we assume a spectral index of $\alpha = 0.75$, considering the synchrotron emission of radio sources which dominates the spectrum of radio sources below 30 GHz (Condon, 1992).

4.3 Dipole estimators

Different types of estimators have been used to measure the cosmic radio dipole, which for the most part yield consistent results. Most commonly used are linear and quadratic estimators (e.g. Singal, 2011; Rubart and Schwarz, 2013; Siewert et al., 2021). Linear estimators essentially sum up all source positions, and therefore, by design, point towards the largest anisotropy in the data. However, the recovered amplitude from linear estimators is inherently biased. Furthermore, because of the sensitivity of linear estimators to anisotropies, any gaps or systematic effects in the data can introduce biases in the estimate of the dipole direction. To avoid biasing the estimator with respect to dipole direction, a mask must be created such that the map remains point symmetric with respect to the observer, or a ‘masking correction’ must be applied (e.g. Singal, 2011; Rubart and Schwarz, 2013). Consequently, missing data features such as the Galactic plane must be mirrored to maintain symmetry,

removing even more data. The quadratic estimator compares expected number counts with a model, providing a chi-squared test of the data with respect to a model of the dipole (e.g. Siewert et al., 2021). The best-fit dipole parameters are then retrieved by minimising χ^2 . Though the cost is the imposition of a dipole model on the data, the estimate is not biased by the ubiquitous spatial gaps in the data of radio surveys.

Both aforementioned estimators are sensitive to anisotropies in the data introduced by systematic effects. Most commonly, these systematic effects influence the sensitivity of the survey in different parts of the sky, meaning the most straightforward solution is to cut out all sources below some flux density. However, given this assessment, we might expect information in terms of the sensitivity of the survey in different parts of the sky to help alleviate these biases. The new generation of radio surveys provides catalogues with a wealth of information, including the local root-mean-square (rms) noise, which we exploit in this work.

To improve sensitivity to the dipole, several attempts have been made to combine different radio catalogues. Both Colin et al. (2017) and Darling (2022) worked on combining different radio surveys, using different techniques to deal with systematic differences between the catalogues. Here, we provide an alternative method to combine catalogues for increased sensitivity to the dipole, while accounting for systematic differences between the catalogues.

4.3.1 Quadratic estimator

To control for differences in pixelation and masking strategies between this and previous works, introducing a new estimator warrants a comparison with known methods of dipole estimation. The quadratic estimator is the closest analogy to the Poisson estimator used here to produce our main results in that it is insensitive to gaps in the data. Its effectiveness and results on multiple large radio surveys are presented in Siewert et al. (2021). The quadratic estimator is based on the Pearson's chi-squared test, minimising

$$\chi^2 = \sum_i \frac{(n_{i,obs} - n_{i,model})^2}{n_{i,model}}, \quad (4.8)$$

where the dipole model is written as

$$n_{i,model} = \mathcal{M}(1 + \vec{d} \cdot \vec{n}_i). \quad (4.9)$$

Here, the dipole amplitude on a given cell is given by the inner product between the dipole vector \vec{d} and the unit vector pointing in the direction of the cell \vec{n}_i , with $\vec{d} \cdot \vec{n}_i = \mathcal{D} \cos \theta_i$. In addition to the dipole vector, the monopole \mathcal{M} is a free parameter, for which the mean value of all cells \bar{n} is a good initial estimate. The χ^2 test is agnostic to the actual distribution of the data, but a dipole model is imposed on the data. This can give rise to misleading results if there are anisotropies in the data—such as those caused by systematic effects—on large enough scales and with sufficiently large amplitude to influence the fit. This can generally be assessed with the reduced χ^2 , which should take a value of around unity if the fit is good.

4.3.2 Poisson estimator

In the dipole estimators used in previous works, no explicit assumption was made as to the shape of the distribution of sources. A Gaussian distribution can be a valid assumption for a source distribution,

although it has an additional degree of freedom compared to Poisson and is less valid if cell counts are low. However, as we do not know a priori how many sources we have, and as we do not commit to a cell size for which we count sources, we choose to assume a Poisson distribution for our cell counts. The Poisson probability density function is given by Equation 4.1, and depends on the mean of the distribution λ , which is equal to the monopole \mathcal{M} in the absence of anisotropies. To account for the effect of the dipole, we introduce a dipole model equivalent to Equation 4.9,

$$\lambda(\vec{d}, \mathcal{M}) = \mathcal{M}(1 + \vec{d} \cdot \vec{\hat{n}}). \quad (4.10)$$

In order to estimate the dipole parameters, we maximise the likelihood, which is given by

$$\mathcal{L}(n|\vec{d}, \mathcal{M}) = \prod_i \frac{\lambda(\vec{d}, \mathcal{M})^{n_i} e^{-\lambda(\vec{d}, \mathcal{M})}}{n_i!}. \quad (4.11)$$

Maximising the likelihood through posterior sampling has the key advantage of immediately yielding the uncertainties on the derived parameter values. This removes the necessity for null-hypothesis simulations as performed by for example Rubart and Schwarz (2013) and Secrest et al. (2021) and Secrest et al. (2022). This estimator was recently used by Dam et al. (2023), who showed that it provides tighter constraints on the dipole than previously used methods.

4.3.3 Poisson-rms estimator

While a survey can be influenced by many different systematic effects, we can assert that the net effect is different sensitivity of the survey at different parts of the sky, leading to anisotropic number counts. Therefore, we assume that all systematic effects that impact source counts in fact influence local noise, thereby causing the source density to vary across the survey. If the survey has sensitivity information in each part of the sky, the impact can be simply modelled by introducing additional variables to the model. So long as a detection threshold is consistently applied to the entire survey, the lower flux density limit will be linearly related to the local rms noise. Consequently, taking into account the dipole and the power law describing number counts, we can model the mean counts in the Poisson estimator as

$$\lambda(\vec{d}, \mathcal{M}, \sigma, x) = \mathcal{M} \left(\frac{\sigma}{\sigma_0} \right)^{-x} (1 + \vec{d} \cdot \vec{\hat{n}}), \quad (4.12)$$

where σ is the rms noise of the cell, x is the power-law index of the flux distribution, and σ_0 is a reference rms value that scales the power law and explicitly ensures λ is dimensionless. The value of σ_0 does not influence any parameters except for the monopole \mathcal{M} , which will take a value closest to the mean cell count \bar{n} when taking σ_0 as equal to the median rms noise over all cells. As the dipole amplitude depends on the the power-law index of the flux distribution near the flux limit, the variation in the rms noise should be small enough that it can be adequately described with a single value. We expect a linear relation between the flux density limit and the local noise, as the detection threshold for most surveys is some multiple of the noise; usually 5σ . Maximising the likelihood given by Equation 4.11 while inserting Equation 4.12 for λ can therefore yield the best-fit dipole and power-law parameters.

4.3.4 Multi-Poisson estimator

Hoping to remove any systematic effects stemming from the incomplete sky coverages of individual radio surveys, Darling (2022) combined the Rapid Australian Square Kilometre Array Pathfinder (ASKAP) Continuum Survey (RACS, McConnell et al., 2020) and the VLA Sky Survey (VLASS, Lacy et al., 2020), finding a dipole that, surprisingly, agrees with the CMB in both amplitude and direction, though with large uncertainties. Secrest et al. (2022) note two inherent problems to this approach. Not only might selecting catalogues at different frequencies select different spectral indices, invalidating the assumption of a common dipole amplitude, but combining the catalogues in such a way ignores systematic effects that can vary between catalogues. Indeed, it is most likely the second factor that plays the most important role, as factors such as observing frequency, array configuration, and calibration can all impact the number counts within a survey in ways that are difficult to predict; this is even more true for independent surveys.

Bearing this in mind, we approach the combination of any two catalogues in a different way. We do not make any attempt to unify the catalogues by matching, smoothing, or creating a common map. Rather, we take both catalogues as independent tracers of the same dipole, allowing the two catalogues to have a different monopole amplitude \mathcal{M} . In the Poisson estimator, we therefore estimate \mathcal{M} separately for each catalogue, turning the likelihood into

$$\begin{aligned} \mathcal{L}(n_1, n_2 | \vec{d}, \mathcal{M}_1, \mathcal{M}_2) &= \prod_i \frac{\lambda(\vec{d}, \mathcal{M}_1)^{n_{1,i}} e^{-\lambda(\vec{d}, \mathcal{M}_1)}}{n_{1,i}!} \\ &\times \prod_j \frac{\lambda(\vec{d}, \mathcal{M}_2)^{n_{2,j}} e^{-\lambda(\vec{d}, \mathcal{M}_2)}}{n_{2,j}!}. \end{aligned} \quad (4.13)$$

Any dipole estimates with this likelihood benefit if the (expected) dipole amplitudes of the two catalogues are similar. However, any differences will be absorbed into the overall error budget by virtue of the sampling algorithm. Once again, the likelihood can be maximised through posterior sampling, yielding the best-fit dipole results as well as the monopoles \mathcal{M}_1 and \mathcal{M}_2 for both surveys.

4.3.5 Priors and injection values

For efficient parameter estimation through posterior sampling, proper priors must be set. While priors on parameters can take on many shapes based on prior knowledge, we take flat priors on all parameters. This only leaves us to define the extent of the probed parameter space, as well as the initial guesses to serve as a starting point for the posterior sampling. In terms of dipole parameters, we separately infer dipole amplitude \mathcal{D} , as well as the right ascension and declination of the dipole direction. We expect the dipole amplitude to take values of around 10^{-2} , but to allow for more variation, the prior on the dipole amplitude is set to $\pi(\mathcal{D}) = \mathcal{U}(0, 1)$. Any point in the sky can represent the dipole direction, and so logically the priors cover the entire sky: $\pi(\text{R.A.}) = \mathcal{U}(0, 360)$ and $\pi(\text{Dec.}) = \mathcal{U}(-90, 90)$. As an initial guess for these parameters, we inject the approximate expected dipole parameters from the CMB: $\mathcal{D} = 4.5 \times 10^{-3}$, R.A. = 168° , Dec. = -7° .

Additionally, the parameters of the distribution of number counts are also estimated. In the basic Poisson case, this is represented by the monopole \mathcal{M} . As the dipole is not expected to meaningfully impact this value, a good initial guess of the monopole is the mean of all cell counts \bar{n} . As the real

Table 4.1: Areas in NVSS masked due to high source density.

Region	RA _◦ _{min}	RA _◦ _{max}	Dec _◦ _{min}	Dec _◦ _{max}	Sky area sq. deg.
1	82.0	90.0	-7.0	-1.0	47.9
2	49.0	52.0	-39.0	-36.0	7.1
3	185.0	189.0	11.0	14.0	11.7
Total					66.7

Table 4.2: Areas in RACS masked due to low source density.

Region	RA _◦ _{min}	RA _◦ _{max}	Dec _◦ _{min}	Dec _◦ _{max}	Sky area sq. deg.
1	357.0	3.0	16.0	22.0	34.0
2	330.0	337.0	16.0	22.0	39.7
3	252.0	261.0	3.0	9.0	53.7
4	184.0	192.0	9.5	15.5	46.8
Total					174.2

monopole is likely close to this value, we choose the prior $\pi(\mathcal{M}) = \mathcal{U}(0, 2\bar{n})$. For the Poisson-rms estimator, both a monopole \mathcal{M} and power-law index x are estimated. Before the estimation, we fit a power law to the cell counts in order to obtain initial estimates \mathcal{M}_{init} and x_{init} , which also function as the initial guesses for these parameters. The initial monopole estimate informs the prior as we use $\pi(\mathcal{M}) = \mathcal{U}(0, 2\mathcal{M}_{init})$. For the power-law index x , a value of around 0.75-1.0 is always expected, and so we take the prior $\pi(x) = \mathcal{U}(0, 3)$.

4.4 Data

Given the estimators introduced in Section 4.3, there are a multitude of available radio catalogues to possibly make use of for a dipole estimate. To maintain the focus of our approach, we use NVSS and RACS, two catalogues that cover the full sky when combined, but have been processed in very different ways given the respective eras in which they were produced. Given the introduction of novel estimators, the NVSS is the logical choice for verification, providing a baseline as the most thoroughly studied catalogue in terms of dipole measurements. The choice of RACS for the second catalogue is straightforward; not only does it complement NVSS in terms of sky coverage, but the inclusion of sensitivity information in the catalogue makes it suitable for testing the Poisson-rms estimator described in Section 4.3.3. The complementary sky coverage of the two catalogues also provides the best testing ground for the Multi-Poisson estimator described in Section 4.3.4.

4.4.1 NVSS

The NVSS (Condon et al., 1998) is one of the most well-studied surveys in terms of dipole measurements, and as such is well suited for verifying novel dipole estimators. It covers the whole sky north of -40° declination, and has a central frequency of 1.4 GHz and an angular resolution of $45''$. The complete catalogue includes the Galactic plane and contains 1,773,484 sources.

An important feature of the NVSS catalogue is that, for observations below a declination of -10° and above a declination 78° , the VLA DnC array configuration was used for observations, while the VLA D configuration was used for the rest of the survey. This affects the number counts at those declinations. The left plot of Figure 4.1 shows the source density of NVSS as a function of declination for different flux density cuts. The impact of the different array configurations can be clearly seen at lower flux densities. At only around 15 mJy, the source density becomes homogeneous, and so for an unbiased dipole analysis we choose to exclude all sources with a flux density below 15 mJy, which is a commonly applied flux density cut (e.g. Singal, 2011; Siewert et al., 2021). Even after this cut, some

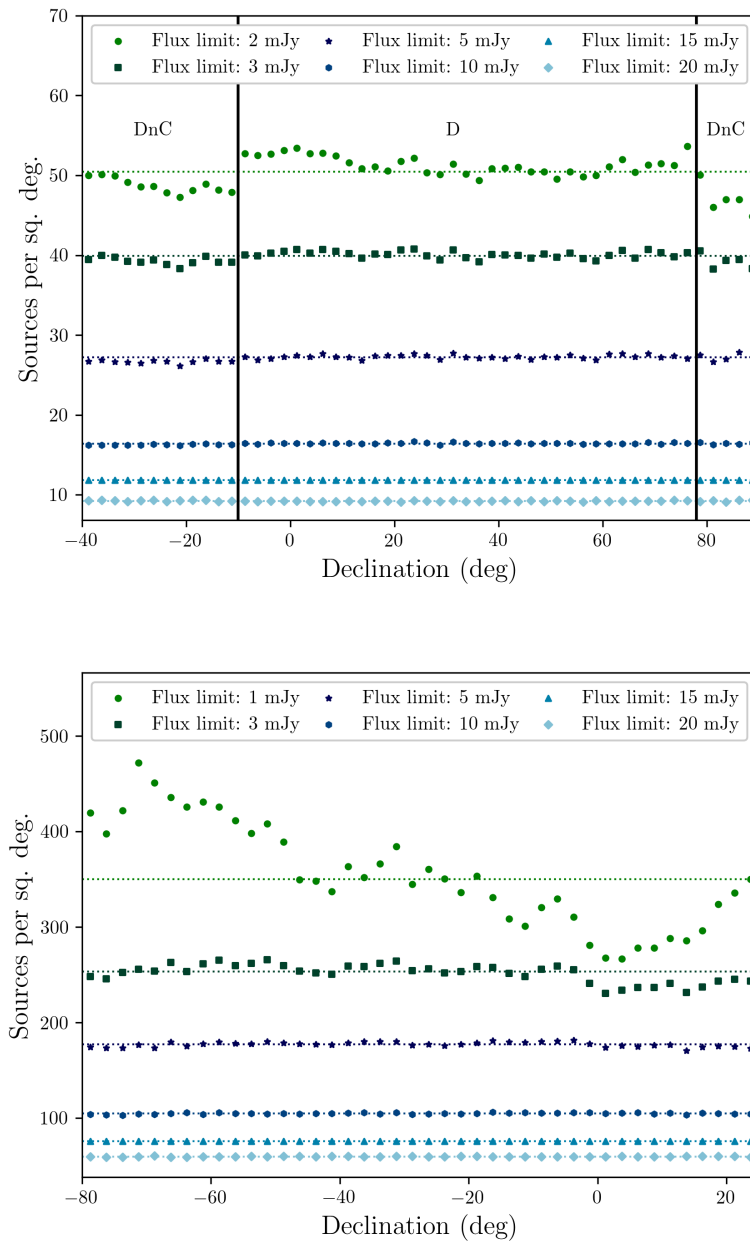


Figure 4.1: Source density of NVSS (top) and RACS (bottom) as a function of declination for different flux density cuts. For the NVSS, we show the boundaries where different array configurations are used. In both cases, the catalogues can be seen to be inhomogeneous at low flux densities.




Figure 4.2: Number counts for NVSS (left) and RACS (right) in equatorial coordinates, in $N_{side} = 32$ HEALPIX maps with masks and a flux density cut of 15 mJy applied.

areas with significantly high source counts are present in the data. We mask these areas as specified in Table 4.1 for the dipole estimate.

With a flux density cut at 15 mJy, we fit a power law to the lower end of the flux distribution of sources and find a power-law index of $x = 0.85$. Additionally taking $\alpha = 0.75$ and taking the velocity from the CMB into account (Equation 4.7), this sets the expectation of the dipole amplitude to $\mathcal{D} = 4.30 \times 10^{-3}$.

4.4.2 RACS

RACS is the first large survey carried out using the Australian Square Kilometre Array Pathfinder (ASKAP), covering the sky south of $+40^\circ$ declination. Observations are carried out with a central frequency of 887.5 MHz and images are smoothed to a common angular resolution of $25''$. The first data release of RACS in Stokes I is described in Hale et al. (2021), and the catalogue used in this work is the RACS catalogue with the Galactic plane removed, containing 2,123,638 sources. Source finding in the images has been done with the Python Blob Detector and Source Finder (PyBDSF, Mohan and Rafferty, 2015), which provides a wealth of information on each source. Most importantly, the rms noise at the position of each source is present in the `noise` column of the catalogue, and we use this in the Poisson-rms estimator approach.

The right plot of Figure 4.1 shows the source density of RACS as a function of declination for different flux density cuts. Even though there is no change in array configuration as with NVSS, there is a clear gradient in source density, where a greater number of sources are detected at lower declinations. Once again, the catalogue becomes homogeneous at flux densities of around 15 mJy, and so we exclude all sources below this flux density, as with NVSS. Even after this cut, some areas with significantly low source counts are present in the data. These areas appear just above the celestial equator. We mask these areas as specified in Table 4.2 for the dipole estimate.

With a flux density cut at 15 mJy and the other masks applied, we fit a power law to the lower end of the flux distribution of sources of RACS and find a power-law index of $x = 0.82$. Taking once again $\alpha = 0.75$, this sets the expectation of the dipole amplitude to $\mathcal{D} = 4.24 \times 10^{-3}$.

4.4.3 Common masks and pixelation

In order to avoid biases to the data, we adopt a masking scheme that uses the same principles for both NVSS and RACS. As mentioned previously, surveys are generally not homogeneous at the lowest flux densities due to variations in the noise, and so we must choose a lower flux density threshold appropriate for the survey. As described above, for both NVSS and RACS we choose a flux density threshold of 15 mJy. To avoid counting Galactic sources or Galactic extended emission, we exclude the Galactic plane. Due to problems with source finding in the Galactic plane, it is already removed in the RACS catalogue, excluding all sources with $|b| < 5^\circ$. For NVSS, overdensities from the Galactic plane extend further out, and we exclude all sources with $|b| < 7^\circ$. Finally, we expect that the brightest sources in a survey push to the limit of the dynamic range, which introduces artefacts around bright sources. Due to differences in source-finding methods between the two catalogues, this increases counts around bright sources in NVSS, but decreases counts around bright sources in RACS. In both cases, the local source density is affected, and therefore we remove all sources within a radius of 0.3° around any source brighter than 2.5 Jy.

After masking data, we use the Hierarchical Equal Area isoLatitude Pixelation (HEALPIX, Górski et al., 2005)¹⁵ scheme to divide the sky into cells of equal size. HEALPIX allows the flexibility of choosing the number of cells over the whole sky, with the base and minimum value being 12 pixels. The resolution parameter N_{side} determines the number of pixels by $N_{pix} = 12 \times N_{side}^2$. We choose two resolutions for our experiment, $N_{side} = 32$ and $N_{side} = 64$, which have pixel sizes of $110'$ and $55'$ on a side, respectively. For measuring number counts, each cell holds the number of sources detected within the confines of the cell. To avoid edge effects resulting from pixels covering data only partially, all pixels that have a neighbouring pixel with zero sources are also set to zero. Finally, all pixels with a value of zero are masked to ensure that these pixels are not taken into account during dipole estimation. The HEALPIX maps of NVSS and RACS with $N_{side} = 32$, with flux density cuts and masks applied, are shown in Figure 4.2. For the Poisson-rms estimator, the local noise is determined per HEALPIX cell by taking the median rms of all the sources within the cell.

4.4.4 Simulations

In addition to the survey data sets of NVSS and RACS, we create a catalogue of simulated sources with a dipole effect to test the validity of the estimators. To do this, we uniformly populate the sky with sources, and assign a rest flux density S_{rest} according to the power law

$$S_{rest} = S_{low} (1 - \mathcal{U})^{-1/x}, \quad (4.14)$$

where \mathcal{U} is a uniform distribution between 0 and 1, S_{low} is the lower flux density limit at which sources are generated, and x the power-law index of the flux density distribution. We transform the rest flux densities and positions of sources by applying relativistic aberration, Doppler shift, and Doppler boost as expressed in Equations 4.4 and 4.5.

We add Gaussian noise to the flux densities of the sources, generating a larger sample of sources and simulating source extraction by only including sources with $S/N > 5$ in the final catalogue. This naturally adds the effect of Eddington bias to the sample, which is expected to be present in real source catalogues. For a realistic distribution, the local noise variation is taken from the RACS $N_{side} = 32$

¹⁵ <http://healpix.sourceforge.net>

Table 4.3: Dipole estimates using the various estimators on the NVSS and RACS catalogues, including a combined estimate using both catalogues.

Survey	Estimator	NSIDE	S_0 (mJy)	N	\mathcal{M} counts/pixel	x	\mathcal{D} ($\times 10^{-2}$)	R.A. (deg)	Dec. (deg)	χ^2/dof
Simulation	Poisson	32	1	4,205,065	489.6 ± 0.3	–	31.7 ± 0.1	202 ± 1	-78.5 ± 0.2	37.3
		32	15	434,247	54.5 ± 0.1	–	1.50 ± 0.25	172 ± 10	-17 ± 13	1.01
		32	50	130,150	16.3 ± 0.1	–	2.06 ± 0.60	178 ± 19	-46^{+17}_{-14}	1.02
	Poisson-rms	32	5σ	4,212,472	518.9 ± 0.3	1.000 ± 0.002	1.58 ± 0.08	169 ± 4	-12 ± 4	1.10
RACS	Quadratic	32	15	451,003	57.4 ± 0.1	–	1.35 ± 0.26	192 ± 12	6 ± 17	1.23
		64	15	458,152	14.74 ± 0.03	–	1.30 ± 0.24	195 ± 11	4 ± 16	1.11
	Poisson	32	15	451,003	56.8 ± 0.1	–	1.41 ± 0.24	193^{+11}_{-10}	4^{+14}_{-15}	1.24
		64	15	458,152	14.19 ± 0.02	–	1.42 ± 0.24	194 ± 10	6 ± 14	1.13
	Poisson-rms	32	5σ	2,035,375	253.3 ± 0.2	0.778 ± 0.003	1.62 ± 0.12	210 ± 4	-12^{+7}_{-6}	2.03
		64	5σ	2,068,204	63.8 ± 0.1	0.738 ± 0.003	2.0 ± 0.2	204 ± 5	-40 ± 4	1.49
NVSS	Quadratic	32	15	345,803	41.0 ± 0.1	–	1.35 ± 0.30	145 ± 13	-5 ± 17	1.17
		64	15	352,862	10.66 ± 0.02	–	1.23 ± 0.29	150 ± 13	-10 ± 17	1.09
	Poisson	32	15	345,803	40.4 ± 0.1	–	1.40 ± 0.29	146 ± 12	-5 ± 15	1.26
		64	15	352,862	10.11 ± 0.02	–	1.38 ± 0.29	151 ± 12	-10 ± 15	1.15
					\mathcal{M}_1 counts/pixel	\mathcal{M}_2 counts/pixel				
NVSS + RACS	Multi-Poisson	32	15	796,806	40.4 ± 0.1	56.7 ± 0.1	1.30 ± 0.18	173 ± 9	-1 ± 11	–
		64	15	811,014	10.11 ± 0.02	14.18 ± 0.02	1.29 ± 0.18	175 ± 8	-2 ± 11	–

rms map, idealising by assigning the rms of a cell to all sources in that cell. Additionally, we simulate false detections by generating sources with the same flux distribution that are not affected by the dipole. These sources consist of 0.3% of the total catalogue, which is the percentage reported for the RACS catalogues (Hale et al., 2021).

All sources are simulated with a spectral index of 0.75, and the power law used to generate the flux distribution has $x = 1$. We apply the dipole effect assuming the direction derived from the CMB dipole, (RA,Dec.) = $(170^\circ, -10^\circ)$, but with an increased velocity of $v = 1107$ km/s to ensure that a sensitivity to the dipole is reached that is similar to that of NVSS and RACS, with similar monopole values. This sets the expectation of the dipole to $\mathcal{D} = 1.5 \times 10^{-2}$.

4.5 Results

Using the described estimators, we estimate the dipole parameters for NVSS and RACS. The results are summarised in Table 4.3, with dipole directions shown in Figure 4.3 and dipole amplitudes shown in Figure 4.4. To estimate the best-fit parameters, for the quadratic estimator, we minimise χ^2 using LMFIT (Newville et al., 2016). The reduced χ^2 -values are reported in Table 4.3. For the Poisson estimators, we use the Bayesian inference library BILBY (Ashton et al., 2019), which provides a convenient and user-friendly environment for parameter estimation. We maximise the likelihood using Markov chain Monte Carlo (MCMC) sampling with EMCEE (Foreman-Mackey et al., 2013). The




Figure 4.3: Best-fit dipole directions for the Poisson estimator of NVSS (blue), RACS (red), RACS with rms power law (green), and NVSS+RACS (purple) compared with the CMB dipole direction (black \odot symbol). Different transparency levels represent 1σ , 2σ , and 3σ uncertainties. In all cases, results from the $N_{side} = 32$ HEALPix map are shown.

scripts used to obtain these results are available on GitHub¹⁶ and an immutable copy is archived in Zenodo (Wagenveld, 2023).

To get an indication of how well the distribution fits a Poisson distribution, Pearson’s χ^2 can also be used as a Poisson dispersion statistic, defined as

$$\chi^2 = \sum_i \frac{(n_i - \bar{n})^2}{\bar{n}}. \quad (4.15)$$

As we use the mean of the distribution instead of a model expectation, the resulting χ^2 value indicates the ratio between the variance and the mean of the distribution¹⁷. For a Poisson distribution, the χ^2 value divided by the number of degrees of freedom (dof), χ^2/dof , should therefore be (close to) unity. In the case of the Poisson-rms estimator, number counts are corrected for the derived power law before calculating χ^2 .

4.5.1 Quadratic and Poisson estimators

Both the quadratic and Poisson estimators are insensitive to gaps in the data, but will be sensitive to inhomogeneous source counts in the data. As such, we perform the flux density cuts on all the

¹⁶ <https://github.com/JonahDW/Bayesian-dipole>

¹⁷ Peebles (1980) uses this measure as a clustering statistic of the large-scale structure, and specifically to define the number of objects per cluster.




Figure 4.4: Best-fit dipole amplitudes with 3σ uncertainties for the Poisson estimator of NVSS (blue), RACS (red), RACS with rms power law (green), and NVSS+RACS (purple), compared with an expected CMB dipole amplitude of $\mathcal{D} = 4.5 \times 10^{-3}$ (black line). Both results from the $N_{side} = 32$ and $N_{side} = 64$ HEALPIX maps are shown.

data before estimating the dipole. As specified in Section 4.4, we choose a flux density cut of 15 mJy for all catalogues in addition to the other masks described above. As the simulated catalogue uses the local noise information from RACS, the same masking is applied there. This leaves $\sim 3.5 \times 10^5$, $\sim 4.5 \times 10^5$, and $\sim 2.2 \times 10^6$ sources for NVSS, RACS, and the simulated catalogue, respectively. The resulting number counts for NVSS and RACS are shown in Figure 4.2. Along with estimating the dipole amplitude and direction, we estimate the monopole \mathcal{M} . For NVSS and RACS, these parameters are estimated using HEALPIX maps of both $N_{side} = 32$ and 64 ; the results of the different cell sizes are shown in Table 4.3.

The best-fit parameters for the simulated data set are shown in Table 4.3, that is, for a low threshold of 1 mJy, the common threshold of 15 mJy, and a high threshold of 50 mJy. As the noise variation of the simulated catalogue is based on RACS, the 15 mJy threshold should be appropriate for obtaining a good estimate of the injected dipole. The low threshold shows that the dominant anisotropy from the RACS noise, which dominates the dipole by three orders of magnitude, is mostly a declination effect, but a smaller effect in right ascension is also observed. With the 15 mJy threshold, the injected values for the dipole are retrieved within the uncertainties. To see how the estimator reacts to a lack of sources, for the 50 mJy threshold, the required number counts are not reached for a 3σ measurement of the dipole amplitude. However, this does not introduce a bias, as values still match the injected values, albeit with large uncertainties.

Comparing results between the quadratic and basic Poisson estimators on NVSS and RACS, values




Figure 4.5: Number counts (left) and median rms (right) for RACS in equatorial coordinates, in $N_{side} = 32$ HEALPix maps, with no flux density cuts applied.

match within the uncertainties for all estimated parameters. For NVSS, the results of the quadratic estimator and Poisson estimator match those of Siewert et al. (2021) for a flux density cut of 15 mJy in terms of dipole amplitude, but the direction is slightly offset ($\Delta\theta \sim 20^\circ$). This is caused by a difference in masking strategy; as shown by Siewert et al. (2021), different masks yield different dipole parameters. The low χ^2 values for the Poisson estimators indicate that a Poisson assumption is in line with the expected distribution of source counts.

As the results in Table 4.3 indicate for both quadratic and Poisson estimators, the results between $N_{side} = 32$ and $N_{side} = 64$ pixel sizes agree with each other within the uncertainties. Furthermore, the dipole amplitudes of NVSS and RACS also agree with each other within the uncertainties, with the dipole amplitude from RACS being slightly higher. The dipole directions between RACS and NVSS are somewhat misaligned ($\Delta\theta \sim 50^\circ$), though both align with the CMB dipole direction within 3σ , that is, at $\Delta\theta \sim 20^\circ$ and $\Delta\theta \sim 30^\circ$ for NVSS and RACS, respectively (see also Figure 4.3). Figure 4.4 shows the amplitudes of the results of the Poisson estimator on NVSS and RACS, including uncertainties. In all cases, the amplitude of the dipole is 3–3.5 times higher than the dipole amplitude expectation from the CMB. For NVSS, the result is at 3.4σ significance and for RACS at a significance of 4.1σ .

4.5.2 Poisson-rms estimator

As described in Section 4.3.3, we aim to account for the variation in source counts across the survey by assuming these are described by the rms noise of the images. For this estimator, we do not apply the flux density cut, and instead fit a power law that relates the rms of a cell to the number counts in that cell. The rms of the survey is not available for NVSS, but is present in the RACS catalogue for each source individually. We obtain the rms of a cell by taking the median rms value of all sources within it. We take the median rms of all cells as the reference rms, which is $\sigma_0 = 0.33$ mJy/beam. The HEALPix maps of source counts and median rms per cell for RACS are shown in Figure 4.5, showing the variation of rms and source counts across the survey. Along with estimating the dipole parameters, the monopole \mathcal{M} and power-law index x are estimated as well. For RACS, these parameters are estimated using HEALPix maps of both $N_{side} = 32$ and $N_{side} = 64$.

The parameters for the simulated data set are estimated and shown in Table 4.3. The noise variation of the simulated catalogue is based on the $N_{side} = 32$ RACS rms map shown in Figure 4.5, which in this case means that the rms map is a perfect representation of the noise in the catalogue. The rms

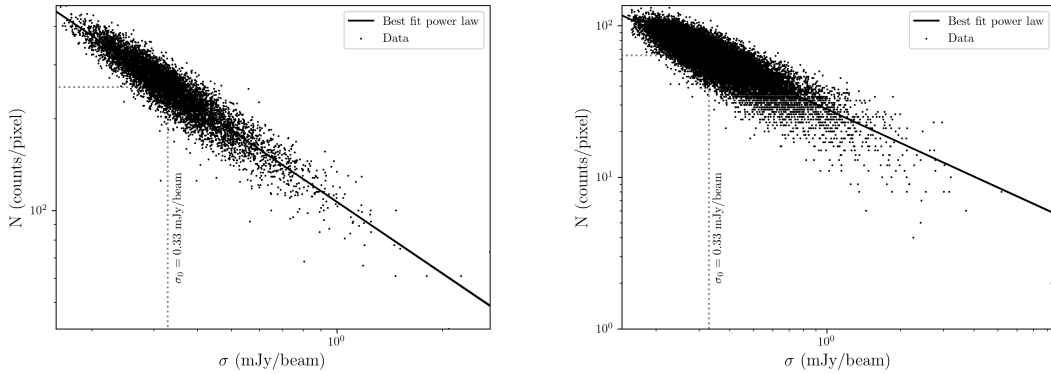


Figure 4.6: Cell counts of the $N_{side} = 32$ (left) and $N_{side} = 64$ (right) HEALPix maps with no flux density cuts applied, as a function of the median rms of the pixels, along with the best-fit power-law model (black solid line). The determined σ_0 is indicated by the dashed vertical line; it intercepts the power law at the best-fit monopole value indicated by the horizontal dashed line.

estimator retrieves the injected dipole parameters, with a much higher significance than the standard Poisson estimator.

The results for RACS are shown in Table 4.3, showing a rather large discrepancy between the two pixel scales, and with respect to other results as well. In both cases, the dipole amplitude is increased with respect to the quadratic and basic Poisson estimators, and the direction is no longer agreeing with the direction of the CMB dipole. The $N_{side} = 32$ map seems to be less affected than the $N_{side} = 64$ map, but in both cases the dipole direction is further away from the CMB dipole direction, with $\Delta\theta \sim 40^\circ$ and $\Delta\theta \sim 45^\circ$ separation for the $N_{side} = 32$ and $N_{side} = 64$ maps, respectively. Especially striking is the recovered dipole direction of the $N_{side} = 64$ map, which is at a declination of -40 deg. This retrieved dipole direction aligns towards the anisotropy retrieved in the simulated data with the 1 mJy flux density threshold. Rather than pointing to an additional systematic effect that is not modelled by the local rms, it is therefore more likely that the median rms noise per cell does not adequately represent the noise variation observed in the catalogue.

To further investigate these results, the power-law fits to the cells are shown in Figure 4.6, indicating that both power laws are a good fit to the distribution. For the $N_{side} = 64$ map, the relation fits less well to the cells with lower number counts, possibly indicating that the power-law assumption breaks down for these cells. One effect that can contribute to this is that, at such low number counts, the median rms will be a less robust measure of the local noise. As is the case for the other RACS results, there is a misalignment in right ascension that is even more pronounced here (see also Figure 4.3). As seen in Figure 4.4, the dipole amplitude is also increased. For the $N_{side} = 32$ map, the dipole amplitude is 3.8 times higher than the CMB expectation with a formal significance of 10σ , and for the $N_{side} = 64$ map the dipole amplitude is 4.7 times higher with a formal significance of 8σ .

4.5.3 Combining RACS and NVSS

Following the procedure laid out in Section 4.3.4, we obtain a combined estimate of the dipole parameters of NVSS and RACS, assuming a common dipole amplitude but independent monopole amplitudes. Although we show that slightly different dipole amplitudes are to be expected between the




Figure 4.7: Dipole amplitudes with 3σ uncertainties compared to the amplitude expected from the CMB from this work and to results from Siewert et al. (2021) and Secret et al. (2022). The results from the different works are separated by horizontal dashed lines, showing the results from this work at the bottom.

catalogues, the degree of this difference depends entirely on the degree to which the inferred dipole is kinetic. Nevertheless, because of the nature of parameter estimation with MCMC, any differences in dipole amplitude between the catalogues will be absorbed into the overall uncertainty of the estimated parameters. In terms of monopole, there is no question as to the difference between the catalogues, as from the estimates of the individual catalogues—even with the same cut in flux density—there are large differences in source density.

Table 4.3 shows the results of the dipole parameters for the combined estimate of NVSS + RACS. Whereas the dipole directions for the individual catalogues were misaligned with the CMB dipole direction, the combined estimate favours a dipole direction that is perfectly aligned ($\Delta\theta = 4^\circ$, see Figure 4.3) with that of the CMB dipole. In line with this finding, the dipole amplitude is reduced with respect to either of the individual catalogues; however, is still in tension with the CMB dipole. The dipole amplitude is three times higher than the CMB expectation, with a significance of 4.8σ for both the $N_{side} = 32$ and $N_{side} = 64$ maps. If we base our belief in a dipole result on its agreement with the CMB dipole in terms of direction, then this is the most significant and reliable result we obtain in this work. It is furthermore the most significant result obtained with radio sources to date, matching the significance of the dipole estimate with WISE AGN from Secret et al. (2022), although less significant than the joint WISE+NVSS result from the same work.

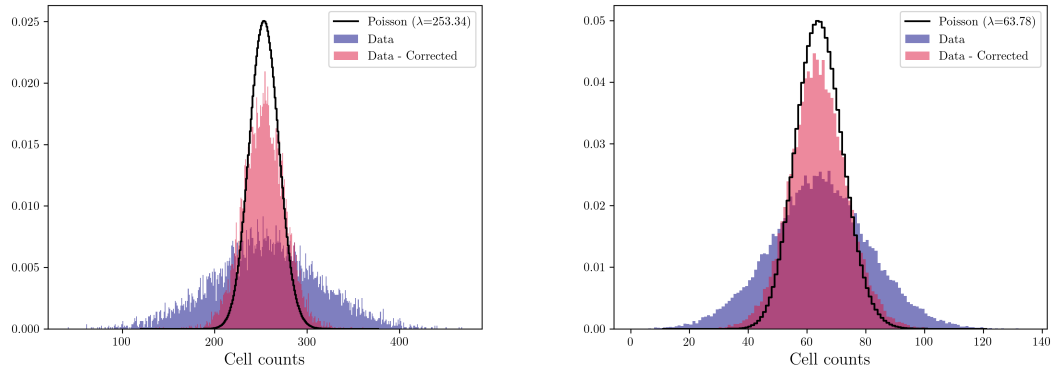


Figure 4.8: Cell count distribution for $N_{side} = 32$ (left) and $N_{side} = 64$ (right) RACS maps without any flux density cuts applied. Raw counts are shown (blue histogram) alongside the counts corrected for the power-law fit (red histogram). A Poisson distribution with a λ equivalent to the estimated monopole amplitude is also shown (black histogram).

4.6 Discussion

The results presented here demonstrate the potential of our introduced estimators and present at the very least an alternative method of making dipole measurements in present and future surveys. Figure 4.7 shows the results from this work compared to the most recent results from various surveys taken from Siewert et al. (2021) and Secrest et al. (2022) in terms of dipole amplitude. The NVSS is consistent across all works, as it is with our results. The results obtained here agree within uncertainties with those from Secrest et al. (2022), with the exception of RACS-rms and the WISE measurements. The same goes for the findings of Siewert et al. (2021), with the exception of the TGSS result. Though the extremely high amplitude of TGSS is attributed to a frequency dependence of the dipole in Siewert et al. (2021), the WISE result from Secrest et al. (2022) does not follow the fitted trend. Secrest et al. (2022) suggest that the TGSS result might deviate due to issues in flux calibration. As we have measurements at different frequencies, we can tentatively check whether results match up with the frequency evolution model of the dipole amplitude from Siewert et al. (2021), which predicts $\mathcal{D} = 2.3 \times 10^{-2}$ at the RACS frequency of 887 MHz. Our RACS result for a flux cut of 15 mJy, which agrees with NVSS, does not follow the trend predicted by the model; however, the 150 mJy TGSS flux density cut made in Siewert et al. (2021) corresponds to a 40 mJy flux cut in RACS (assuming $\alpha = 0.75$). Applying this flux density cut using the $N_{side} = 32$ RACS map, the inferred dipole direction shifts by $\Delta\theta = 8.8^\circ$, and the dipole amplitude increases to $\mathcal{D} = (1.94 \pm 0.37) \times 10^{-2}$. As such, our results cannot rule out the frequency dependence predicted by Siewert et al. (2021), but the obtained results from WISE AGN (Secrest et al., 2021; Secrest et al., 2022; Dam et al., 2023) provide a strong argument against it. Though our results are consistent with the literature, as with many works concerning the dipole, their validity and that of the methods require further examination.

4.6.1 The Poisson solution

Though we show results here that are both internally consistent and consistent with other dipole estimates, the choice of a Poisson estimator might seem like an unnecessary constraint on the data;

after all, the quadratic estimator shows an adequate performance and does not suffer any loss in precision compared to the Poisson estimator. Table 4.3 lists the χ^2 values for the obtained results, defined by Equations 4.8 and 4.15 for the quadratic and Poisson estimators, respectively. As the quadratic estimator is minimised for a Gaussian distribution with mean equal to the variance, the quadratic and basic Poisson estimators are expected to provide similar values. This is indeed the case for the results in Table 4.3, both in the estimated parameters and χ^2 values.

The value of the Poisson assumption becomes readily apparent when extending the parameter space, as we do when taking into account the rms power-law relation. The main feature of a Poisson distribution is that one parameter is necessary to describe it, λ , which is both the mean and variance of the distribution. This is a strict requirement on a distribution, allowing more freedom in other parameters which would otherwise be degenerate with the parameters of the distribution. This means that fitting the rms power law does not work with a quadratic estimator for example; indeed, this latter, though minimised by a distribution with mean equal to the variance, still allows for a wider Gaussian distribution. As seen in Figure 4.8, the distribution of number counts without any flux density cut applied resembles a Gaussian distribution, which is much wider than a Poisson distribution with the same mean. However, the quadratic estimator does allow such a wide distribution, and therefore will not converge on a solution that transforms this distribution to a Poisson distribution. Herein lies the power of the Poisson estimator, which makes modelling and fitting of systematic effects in the data a viable alternative to cutting and masking data. Nevertheless, one drawback is that it is imposing a Poisson distribution on the data, which can lead to spurious results if improperly applied.

Table 4.3 lists the χ^2/dof values of the Poisson rms estimator after correction for the derived power law. The difference in distributions can be appreciated in Figure 4.8, which shows the distributions of the cell counts of RACS without any flux density cut applied, along with the same distribution corrected for the rms power law that has been fit to the data. The uncorrected counts have a much wider distribution, which is clearly not Poisson, with $\chi^2/dof = 13.28$ for the $N_{side} = 32$ map and $\chi^2/dof = 4.49$ for the $N_{side} = 64$ map. The corrected counts resemble a Poisson distribution more closely, with $\chi^2/dof = 2.03$ for the $N_{side} = 32$ map and $\chi^2/dof = 1.49$ for the $N_{side} = 64$ map, but χ^2/dof values indicate variance is still too large for a Poisson distribution, signifying that some residual effect has not been modelled by the estimator.

As such, the performance of the Poisson-rms estimator still leaves some questions to be answered. The assumption that source counts are related to sensitivity via a power law might carry a flaw, though there can be a number of possible reasons for this: (i) the median rms is not the best representation of the sensitivity of the survey in a given cell; (ii) the sensitivity only properly represents source counts down to some limit; and (iii) not all systematic effects equally impact source counts as well as sensitivity. These factors require further examination in the future, but remarkable already are the results when compared to the other RACS results. It is clear that this Poisson-rms estimator shows promise even in its basic form, and can be used as an additional test of the data for any survey that has information on the local rms. Furthermore, due to its flexibility, additional effects once characterised can easily be modelled and taken into account by the estimator.

4.6.2 Residual anisotropies in the data

In dipole measurements and other statistical studies that require large amounts of data to retrieve a statistically significant measure, it can be difficult to visually assess whether any one fit adequately describes the data. After all, we impose a model on the data to which the fit is restricted. For a




Figure 4.9: Dipole amplitude of NVSS as a function of angular distance to the dipole direction, assuming a best-fit dipole direction from the basic Poisson estimator, for both $N_{side} = 32$ (blue) and $N_{side} = 64$ (red) maps. Alongside the data, the corresponding models are plotted (solid lines), as well as the expected model from the CMB dipole (black dashed line).

rudimentary visual verification of whether or not the data follow the expected relations, we employ the hemisphere method used by Singal (2021). This method assumes that the direction of the dipole is already known, leaving the dipole amplitude as a function of angular distance from the dipole direction, $\mathcal{D}_\theta = \mathcal{D} \cos \theta$, as the only free parameter. To reach statistically significant number counts, the sky is divided into two hemispheres: hemisphere N_1 with all sources between θ and $\theta + \pi/2$, and hemisphere N_2 with all sources between $\theta + \pi/2$ and $\theta + \pi$. The dipole amplitude as a function of θ is then written as

$$\mathcal{D}_\theta = \frac{N_1(\theta) - N_2(\theta)}{\frac{1}{2}[N_1(\theta) + N_2(\theta)]}. \quad (4.16)$$

We determine and plot the hemisphere results for NVSS and RACS assuming the results obtained from the Poisson estimators for the individual catalogues. The hemisphere relation for NVSS is presented in Figure 4.9, which shows the data following the expected dipole curve except for the hemispheres closest to the dipole direction; these data reveal an increased anisotropy. This is more pronounced in the $N_{side} = 64$ map, where both $\theta = 0^\circ$ and $\theta = 10^\circ$ hemispheres show significantly increased counts compared to the expectation from the dipole model. In the $N_{side} = 32$ map, only the $\theta = 0^\circ$ hemisphere shows increased counts, with all other points following the dipole model within uncertainties. This points to a residual anisotropy left in the data that has not influenced the overall fit.

For RACS, the hemisphere relations for both the basic Poisson estimator and the Poisson-rms estimator are shown in Figure 4.10. There is a residual anisotropy in RACS at $40^\circ - 60^\circ$ from the

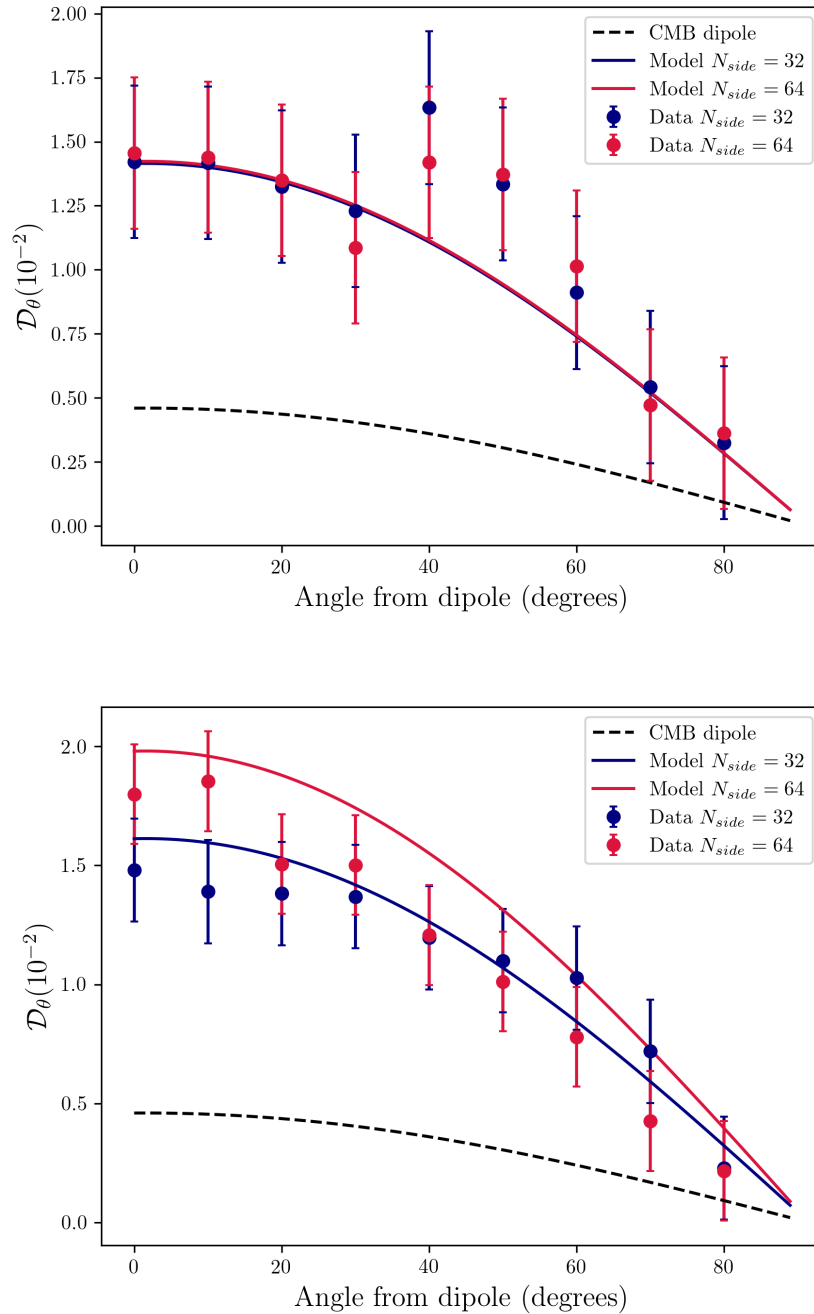


Figure 4.10: Dipole amplitude of RACS as a function of angular distance to the dipole direction, assuming a best-fit dipole direction from the basic Poisson estimator (top) and Poisson rms estimator (bottom). Results for both $N_{side} = 32$ (blue) and $N_{side} = 64$ maps are shown. Alongside the data, the corresponding models are shown (solid lines), as well as the expected model from the CMB dipole (black dashed line).

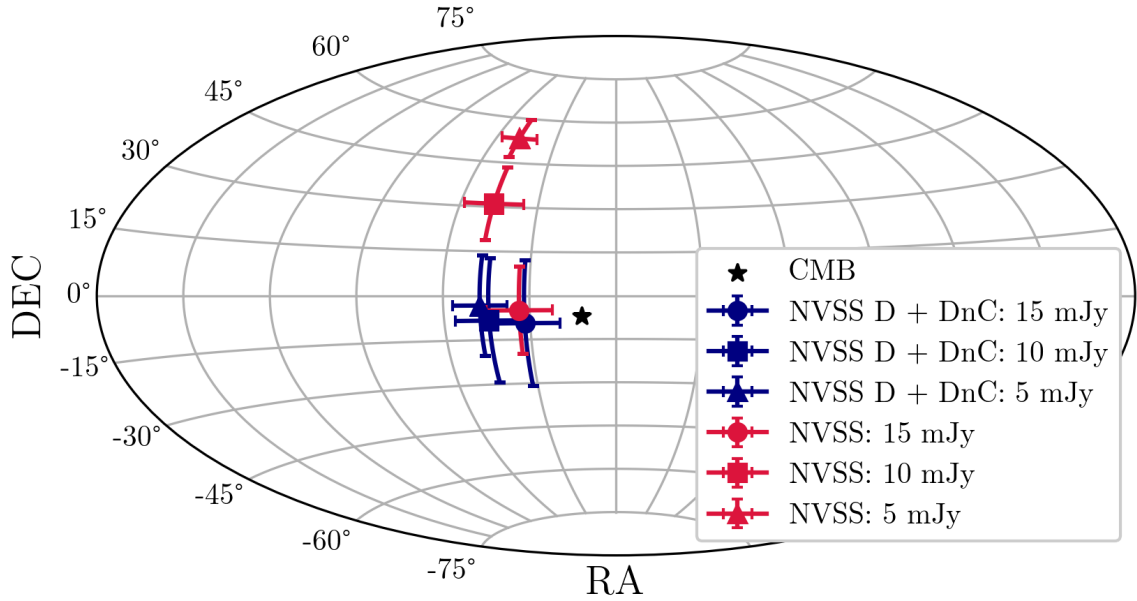


Figure 4.11: Best-fit dipole directions with 1σ uncertainties for the complete NVSS (red) and split NVSS D + DnC (blue) catalogues, for lower flux density thresholds of 5, 10, and 15 mJy. The CMB dipole direction is indicated with a black star.

dipole direction for both $N_{side} = 32$ and $N_{side} = 64$ maps that stands out immediately, especially in the results for the basic Poisson estimator. In the case of the basic Poisson estimator, as with NVSS, the fit is unaffected. However, this anisotropy might have had a significant impact on the Poisson-rms estimator for the $N_{side} = 64$ map, as the dipole direction estimated from that map is 47° offset from the direction of the basic Poisson estimator, coinciding perfectly with the anisotropy seen at that angle. Indeed, the data for both $N_{side} = 32$ and $N_{side} = 64$ maps for the Poisson-rms estimate agree well with the exception of the points closest to the dipole direction, which in the case of the $N_{side} = 64$ map are dominating the fit.

Finally, we investigate the possibility that residual systematic effects are present due to Galactic synchrotron. Secrest et al. (2022) use the de-stripped and source-subtracted Haslam et al. (1982) 408 MHz all-sky map from Remazeilles et al. (2015) to mask pixels bright in Galactic synchrotron. To investigate if our results are impacted by Galactic synchrotron, we cross-correlate our number count maps with the Remazeilles et al. (2015) map. In all cases, no significant correlation is found ($|\rho| < 0.05$), showing that Galactic synchrotron is not present as a residual systematic effect in the data.

4.6.3 Combining and splitting catalogues

As shown in Section 4.5.3, combining catalogues as independent tracers of the dipole can provide a more robust measurement of the dipole, and in the case of NVSS and RACS, result in a dipole that matches the direction of the CMB dipole remarkably well, with a dipole amplitude of 2.5 times the CMB expectation. The justification for this approach is that radio data are a complex product and

Table 4.4: Dipole estimates for NVSS using different flux density cuts, separating D and DnC configurations.

Catalogue	S_0 (mJy)	N	\mathcal{M}_1 counts/pixel	\mathcal{M}_2	\mathcal{D} ($\times 10^{-2}$)	R.A. (deg)	Dec. (deg)
NVSS	15	345,803	40.4 ± 0.1	–	1.40 ± 0.29	146 ± 12	-5 ± 15
	10	480,446	56.1 ± 0.1	–	1.38 ± 0.26	133 ± 11	31 ± 12
	5	795,135	92.6 ± 0.1	–	1.92 ± 0.23	131 ± 9	53 ± 6
NVSS D + NVSS DnC	15	333,046	40.5 ± 0.1	40.4 ± 0.1	1.44 ± 0.31	148 ± 12	-9 ± 21
	10	463,398	56.4 ± 0.1	55.7 ± 0.2	1.23 ± 0.27	136 ± 12	-8_{-20}^{+22}
	5	767,832	93.5 ± 0.2	91.3 ± 0.2	1.18 ± 0.20	133 ± 9	-3_{-16}^{+18}

are sufficiently difficult to homogenise over a full survey internally, not to mention between surveys. Field of view, frequency, array configuration, calibration, imaging and source finding are all factors to consider when assessing the source counts in a given survey. For a key example of how these factors can influence source counts, we need to look no further than NVSS, which has been observed with two different array configurations at certain declination ranges. This is a systematic effect that produces different source densities depending on the configuration, something that can be plainly seen in the left plot of Figure 4.1.

To demonstrate the potential of the Multi-Poisson estimator beyond combining independent catalogues, we repeated the dipole estimates for NVSS, lowering the minimum flux density in several steps. We then split the NVSS D and DnC configurations into separate catalogues, and repeated the experiment. The results are shown in Table 4.4 and Figure 4.11, showing the effect of splitting the configuration on the dipole estimates. Immediately, we can see that, though the dipole estimate in both cases remains consistent in terms of right ascension, the separation of the catalogue produces wildly different results for flux density cuts below 15 mJy, both in terms of dipole amplitude and declination of the dipole direction. This result is expected to some degree, as the difference in source density between the D and DnC configurations is expected to produce an anisotropy in the north–south direction, which is largely alleviated (although not entirely) with the split in configurations. This is not only reflected in the declination of the dipole direction, but in the dipole amplitude as well, which is seen to decrease with more samples in the case of split configurations. Though the anisotropy in the declination is alleviated by this, another anisotropy in right ascension seems to start dominating at lower flux densities, dragging the dipole direction 43° from the CMB dipole direction.

Though it seems that the different NVSS configurations produce an anisotropy that is on a similar level to other anisotropies related to incompleteness of the catalogue, and cannot therefore be reliably used to completely account for the systematic effects in the survey that appear when employing lower flux density thresholds, these results show that were such systematic effects to dominate the catalogue, restructuring the problem to consider these as multiple independent catalogues can produce sensible results. Our analysis is furthermore a useful test of the chosen flux density threshold, as for an appropriately chosen flux density threshold the results between the full catalogue and split catalogue should be consistent. The results here show that a flux density threshold of 15 mJy is indeed appropriate for the NVSS. Such an approach might become particularly relevant given that VLASS

also uses different array configurations, which can be taken into account when estimating a dipole in the same manner as we have done for NVSS. Even for RACS, a prominent feature is a dependence of source counts on declination; though the exact mechanism is unclear, one possibility is related to the point spread function, as the UV-coverage of the array evolves with declination.

4.6.4 Combining catalogues and cosmological considerations

The combined dipole estimate of RACS and NVSS makes a compelling case for combining more probes of the cosmic dipole to increase sensitivity. The approach used here ostensibly carries less caveats than previous works, foregoing source matching, frequency scaling, subsampling, and weighting schemes, which can all introduce additional uncertainties. This carries with it a reduction in formal uncertainty, though there remain some factors regarding the nature of the dipole that can limit the approach. As the nature of the excess amplitude of the radio dipole is currently unknown, the approach of combining catalogues, even if done perfectly, carries an additional uncertainty. In Section 4.4, the expected kinematic dipole amplitudes for both NVSS and RACS are computed and are found to be nearly identical. This in large part justifies the obtention of a combined estimate of the catalogues; however, if we were to combine catalogues where the expected dipole amplitudes differed (e.g. when combining catalogues from multiple wavelengths), the approach would not be able to produce a reliable result without knowing the nature of the radio dipole. As it stands, we have a kinematic expectation of the radio dipole derived from the CMB. Given the velocity of the observer, this dipole is determined by the spectral index and flux distribution of the ensemble of sources. However, because of the measurement method used here, the excess dipole has an unknown origin, be it either kinetic or due to some entirely different effect.

Given the results obtained so far in this work and the literature, one could also assume an effect that somehow boosts the observed dipole with respect to the CMB dipole, or an additional anisotropy that is simply added to the kinematic dipole. Therefore, should we wish to combine for example the sample of WISE AGN, which has an expected dipole amplitude of $\mathcal{D} = 7.3 \times 10^{-3}$ (Secrest et al., 2022), with for example NVSS, a combined estimate would have to assume one of these models. Secrest et al. (2022) find that after removing the CMB dipole, assuming it is purely kinematic, the residual dipoles between NVSS and WISE agree with each other, favouring the interpretation of an intrinsic dipole anisotropy in the CMB rest frame. The results we obtain for NVSS and RACS also support this interpretation, the residual dipole amplitudes being $\mathcal{D} = (0.97 \pm 0.30) \times 10^{-2}$ and $\mathcal{D} = (0.99 \pm 0.24) \times 10^{-2}$, respectively. However, the expected dipoles of the catalogues are too similar to rule out other interpretations.

Furthermore, depending on which of the models presented above is the most accurate, survey design can have a profound impact on the measured cosmic radio dipole. The largest impact will be in the detected source populations and their redshift distributions. Naturally, going to optical or infrared wavelengths will yield different source populations that possibly trace the dipole differently, but even amongst radio surveys the detected source population will depend on the survey details. Radio surveys must be designed with a balance of depth and sky coverage, and so we imagine a scenario where the number of sources detected will stay constant over the survey due to this balance, leaving the significance of a dipole measurement unchanged. A shallow but large-sky-coverage radio survey will mostly detect AGN with a peaked redshift distribution, whereas a deep radio survey with limited sky coverage will probe most of the AGN population over all redshifts, as well as star-forming galaxies, which have a different redshift distribution from AGN altogether. Most surveys that have been used

for dipole estimates fall into the first category, as proper coverage along the dipole axis is necessary. However, a survey falling into the second category would have the potential to differentiate between the possible models of the radio dipole we have laid out. The MeerKAT Absorption Line Survey (Gupta et al., 2016), consisting of sparsely spaced deep pointings homogeneously distributed across the sky, provides a good candidate for a survey falling into this second category, and therefore might provide more insight into the processes driving the anomalous amplitude of the radio dipole.

4.7 Conclusion

In this work, we present a set of novel Bayesian estimators for the purpose of measuring the cosmic radio dipole with the NVSS and RACS catalogues. Based on the assumption that counts-in-cell of radio sources follow a Poisson distribution, we construct estimators for the cosmic radio dipole based on Poisson statistics. To provide a means of for comparison, we include a quadratic estimator, which has been used in a number of previous dipole studies. We furthermore construct two extensions of the basic Poisson estimator to attempt to account for systematic effects in the respective catalogues. Firstly, we consider that if sensitivity information is present in the catalogue, this can be directly related to the local number density, assuming that systematic effects merely modify the local sensitivity of the catalogue. The local sensitivity and number counts are assumed to be related by a power law, the parameters of which can be estimated. We extend the Poisson estimator to address this. Secondly, we construct an extension to the Poisson estimator that can be given multiple separate catalogues—assuming that the catalogues trace the same dipole—and produces a combined estimate.

We obtain best-fit parameters for the cosmic radio dipole using χ^2 minimisation for the quadratic estimator and using maximum likelihood estimation for the Poisson estimators. To discretise the sky, we use HEALPIX, producing maps with both $N_{side} = 32$ and $N_{side} = 64$. We verify that the quadratic estimator and basic Poisson estimator yield similar results, and that furthermore results between the pixel scales are consistent. We use the Poisson-rms estimator on RACS while not using any cut in flux density to estimate the dipole parameters along with the parameters for the rms power law. The increased number counts greatly increase the precision of the estimate, but the results somewhat diverge from the dipole estimates produced by the basic estimators. Whether this difference is a genuine product of the data or a flaw in (assumptions of) the estimator is not perfectly understood, but the initial results are still promising given that the entire catalogue of sources is used. We finally use a Poisson estimator for multiple catalogues on NVSS and RACS and obtain a dipole estimate that perfectly aligns with the CMB dipole in terms of direction, but has an amplitude that is three times as large with a significance of 4.8σ . Given the dipole estimates obtained from the individual catalogues, this result is in line with expectations for a combination of the two catalogues, and can therefore be seen as the most reliable and significant result obtained here.

We explore the possibility of splitting up a catalogue and using the Poisson estimator for multiple catalogues to estimate the dipole as if on two independent catalogues. We use this method on the NVSS, which has been observed with two different array configurations, introducing an artificial north–south anisotropy in the catalogue. We treat these array configurations as separate catalogues and repeat the dipole estimate, and go down to lower flux density limits than with the basic Poisson estimator. We see that while using the whole NVSS, a north–south anisotropy starts to dominate the estimate at lower flux densities; separating the configuration largely mitigates this effect. As a result, this allows us to lower the flux density cut, increasing number counts and thus increasing the

significance of the dipole estimates. This approach may work well on catalogues such as VLASS, which also uses different array configurations in different parts of the sky. The presented estimator may provide the potential to combine a larger variety of catalogues, but the extent to which this can be done depends in large part on the nature of the excess dipole. With an increasing array of probes at various wavelengths, sky coverage, and depth, reaching the necessary sensitivity to detect the dipole, the nature of this dipole may well soon be discovered, given the different populations of sources that can be probed.

Revenge of the systematics

With 391 pointings observed in the L band, the full MALS survey boasts more than a thousand square degrees of deep continuum sky and hundreds of thousands of sources, making it suitable for a dipole measurement. To fully utilise the survey, we have characterised systematic effects present in the pointings with a deep analysis of the calibration, imaging, source extraction, and cataloguing of ten pointings of MALS in Chapter 3. We have furthermore defined a set of Bayesian estimators that will work on the sparse structure of the MALS sky coverage in Chapter 4. Using this, we aim to measure the cosmic radio dipole on the full set of 391 pointings observed in L band.

This chapter is organised as follows. In Section 5.1 we describe the MALS data, and the systematic variation in source density as a function of declination present in Section 5.2. In Section 5.3, we show the differential number counts of the entire catalogue and compare them to other surveys. In Section 5.4 we describe the steps we take to prepare the MALS data for a dipole measurement, including the dipole estimators we use for the measurement and the creation of simulated data sets to test the dipole estimators for biases. Results of the dipole estimates using the MALS data are described in Section 5.5. In Section 5.6 we discuss potential causes of the anisotropy and similar effects appearing in other catalogues. We summarise and conclude in Section 5.7.

5.1 MALS data

The sky distribution of the 391 pointings of MALS observed in L band is shown in Figure 5.1. These pointings have been observed between April 1, 2020 and January 18, 2021. Observations are carried out in 32K mode, splitting the total bandwidth of 856 MHz into 32,786 spectral channels, with a channel width of 26.123 kHz. Typical observation runs include three target pointings, which are observed consecutively for 20 minutes at a time, going back to the first target after the final target is observed. Repeating this three times yields a total of 60 minutes of integration time for each target. Shortly before and after each target observation a nearby gain calibrator is observed for a few minute. Flux density calibrators are observed at the beginning, middle, and end of each observation run for 10 minutes. The flux density calibrators used by MALS are 3C 286, 3C 138, PKS 1939-638, and PKS 0408-658.

More details of MALS observations and calibration using the Automated Radio Telescope Imaging Pipeline (ARTIP, Gupta et al., 2021) are described in Chapter 3.2 and Deka et al. (2023). Each of




Figure 5.1: Sky distribution of 391 MALS pointings that have been observed and processed in L band, in equatorial coordinates. The colouring indicates the number of sources in each pointing, with the dotted lines indicating the galactic latitude range $|b| = 10^\circ$.

these pointings has been calibrated, self-calibrated and imaged in ARTIP. For wideband continuum imaging, the data are averaged over 32 channels and are divided into 15 spectral windows, leading to 64 channels per spectral window. After flagging of frequencies with known radio frequency interference as well as the edges of the band, the total bandwidth is reduced to 802.5 MHz. The procedure of determining the wideband primary beam is the same as described in Chapter 3, however for the full set of pointings `katbeam` primary beams are used instead of the holographic measurements. Furthermore, no primary beam cut off is applied, such that a full 6000 x 6000 pixel image is retained even after primary beam correction. With a pixel size of $2''$, the images are 3.3° on a side. However, as the sensitivity decreases while direction dependent effects increase the number of artefacts further out from the primary beam, we limit the inclusion of sources for a dipole measurement out to a radius of 1.15° , which corresponds to the 5% level of the wideband primary beam.

As described in Chapter 3, we evaluate the logs produced during processing by ARTIP to assess and look for errors and warnings in the pipeline. An important quality to assess is the flux density scale, as systematic errors can easily carry on through to the dipole estimate. The flux density scale is set by the flux density calibrator and subsequently applied to the gain calibrator, so a good check is to compare the flux density of the gain calibrators determined during calibration with a reference catalogue. Figure 5.2 shows precisely this, comparing the flux density of gain calibrators measured




Figure 5.2: Flux density of gain calibrators as determined during calibration compared to the reference flux density values of the same gain calibrators from (Taylor and Legodi, 2021). Most gain calibrators have been used for multiple targets, but on each occasion has a slightly different measured flux density.

during calibration with the flux densities from the MeerKAT reference catalogue (Taylor and Legodi, 2021). The flux densities match well, with a median flux density ratio of 1.03 ± 0.07 . This matches the result obtained from the first ten pointings, though the median ratio is closer to unity, the uncertainty remains the same, indicating that this is likely the intrinsic variance in the flux density scale introduced in the processing pipeline.

5.1.1 Noise properties

Source extraction in the images has been carried out using the Python Blob Detection and Source Finder (PyBDSF, Mohan and Rafferty, 2015). The PyBDSF setup used for the wideband MALS images matches the one used for MALS DR1 described in Deka et al. (2023). This slightly differs from the setup used in Chapter 3 by letting PyBDSF determine the size of the root mean square (rms) smoothing box around bright sources, rather than setting the size manually.

One of the products of PyBDSF source finding are rms maps of each image, which can be used to assess the noise properties of the pointings. The rms maps for these MALS images cover the full $3.3^\circ \times 3.3^\circ$ image size, but as we are only using sources up to 1.15° from the pointing centre, we cut off the rms maps at that distance. In Chapter 3, we defined the rms coverage of the image as the cumulative distribution of pixel values in the rms map. Most pointings show very similar rms coverage curves, the only difference often being an offset in overall noise level. To describe the offset between the pointings, we defined σ_{20} as the rms value at 20% of the cumulative rms coverage. Figure 5.3 shows the rms coverages of all 391 pointings, normalised by their σ_{20} values, showing that




Figure 5.3: RMS coverage of all 391 pointings normalised by their σ_{20} values, along with the median coverage (solid black curve) of all pointings.




Figure 5.4: Median rms map of all 391 pointings as a function of σ_{20} , showing increased noise in the pointing centre and towards the edges of the image. The contour lines show the levels $(1, 2, 4) \times \sigma_{20}$ in the image. For the pointing with the lowest σ_{20} , J2339-5523, the range $1 - 5\sigma_{20}$ corresponds to 12-60 $\mu\text{Jy}/\text{beam}$. For the pointing with the highest σ_{20} , J1244-0446, the range $1 - 5\sigma_{20}$ corresponds to 0.4-2 mJy/beam.




Figure 5.5: Demerit scores for all MALS pointings computed from the catalogues produced by PyBDSF (left) and computed from NVSS and SUMSS (right), as a function of σ_{20} of the pointings. There is a clear correlation between σ_{20} and the demerit score, with a correlation of $\rho = 0.45$ in both cases.

the pointings all follow the same curve, and their differences can be adequately described by their σ_{20} values. The highest spread in the rms coverage curves occurs around $3\sigma_{20}$, with a standard deviation of 5% of the total coverage.

We obtain a smoothed rms map by median stacking the PyBDSF rms maps, normalised by their σ_{20} values, of all 391 pointings. The median rms map is shown in Figure 5.4, in terms of σ_{20} , showing the overall structure of the MALS pointings. The presence of a bright central source in each pointing increases the noise in the centre of the image, and towards the edges of the image the noise also increases following the primary beam response. Differences between individual pointings, such as bright off-axis sources, are washed away in the median stacked image. As would be expected, this median rms map is much smoother compared to the one created with only ten pointings, shown in Figure 3.3. The noise structure associated with the central source is slightly elongated in the north-south direction, which is likely a result of the (averaged) shape of the synthesised beam.

For an alternative measure of image quality we use the demerit score introduced by (Mauch et al., 2020), which computes the impact of bright sources in the pointing to errors and noise in the image. Figure 5.5 shows the demerit score of the pointings computed with Equation 3.8, compared to the σ_{20} values of the pointings. The scores are computed both using the source contributions of bright ($S > 100$ mJy) sources. One set of scores is calculated from external catalogues, specifically the NRAO VLA Sky Survey (NVSS, Condon et al., 1998) and the Sydney University Molonglo Sky Survey (SUMSS, Mauch et al., 2003). Demerit scores are also computed using the sources catalogues of the pointings produced by PyBDSF. There is a clear correlation ($\rho = 0.45$) between σ_{20} and the demerit scores, though there is still a significant amount of variance. This shows that overall, the demerit score is a viable predictor of image quality.

5.1.2 Source catalogues

As some direction dependent effects are not taken into account during imaging, dynamic range of the image decreases further away from the pointing centre, showing an increased number of artefacts around bright sources that are mistakenly identified as sources by the automatic source extraction




Figure 5.6: Number of sources as a function of σ_{20} in each pointing, along with the best fit power law to the data, excluding all pointings with $\sigma_{20} > 45 \mu\text{Jy}/\text{beam}$, indicated by the red points.

from PyBDSF. We therefore expand the artefact flagging method from Chapter 3 to take this effect into account, and identify all sources in the image for which

$$\frac{S}{N} \cdot \rho^2 > 100, \quad (5.1)$$

with ρ being the distance from the pointing centre, in degrees. Here we define the S/N as the ratio of the peak flux density of the source to the local rms noise. Around these bright sources, we flag all sources that are located within 10 times the major axis of the synthesised beam that have a flux density of less than 5% of the bright source. With distance from the pointing centre limited to 1.15° , each pointing covers a sky area of 4.15 deg^2 . With artefacts flagged, the combined source catalogue of all 391 pointings contains 796,916 sources, and the total sky coverage of the survey is 1623 deg^2 . With a lower sensitivity limit of $10 \mu\text{Jy}/\text{beam}$, the faintest sources detected in the survey have a flux density of $50 \mu\text{Jy}$. The number of sources in each pointing is shown in Figure 5.1, showing that on average source counts are higher towards the south pole.

Figure 5.6 shows the number of sources in each pointing as a function of σ_{20} . We see that a number of pointings have very low source counts ($\lesssim 1000$ sources per pointing), usually because of strong sources being either in the pointing centre or elsewhere in the field, which can additionally lead to an increased fraction of false detections. As the quality of these pointings is suspect, we identify all pointings with $\sigma_{20} > 45 \mu\text{Jy}/\text{beam}$, a total of 41 pointings, as low quality pointings. Because of their low number counts, these pointings only have a combined 3.4×10^4 sources. In addition to this, 41 pointings have an associated galaxy cluster in the Atacama Cosmology Telescope (ACT) catalogue (Hilton et al., 2021). There is however no significant difference in source density between

these pointings and pointings without an associated cluster.

5.2 A declination systematic

Figure 5.1 shows the location and number of sources for all MALS pointings. It is apparent that there is a significant variation in source density as a function of declination. Figure 5.7 displays the number of sources per pointing as a function of declination for different flux density cuts, both for each pointing separately (top plot) and averaged in declination bins (bottom plot). The binned counts show that though the effect is most apparent when no flux density cuts are applied, it persists at higher flux density cuts. Though the noise structure of the survey follows a similar trend in declination and is likely caused by the same effect, the effect persists well above the completeness limit of the survey. This systematic variation in source density can dominate the dipole signal if it can not be accounted for. An anisotropy in declination is strongly suggests an observational effect, as the coordinate system is only significant for earth based observations. Given that the effect is biasing results in terms of declination (right ascension seems only minimally affected), the most likely effects have to do with the projection of the array configuration.

The (u, v) -coverage for a specific observation describes the baselines of the array as seen from the observing target, over the observation time. The biggest impact on (u, v) -coverage is the elevation of the target, θ_{el} , as this will determine the projection of the baselines onto the (u, v) plane. The elevation of a target will depend on the time of observation, which is not publicly available information in general, however the maximum elevation a target field can reach is determined by its declination

$$\theta_{el,max} = 90^\circ - |\delta_{array} - \delta_{target}|. \quad (5.2)$$

The elevation angle describes then the angle between the horizon and the target at the location of the array, and thus also the angle which the array is rotated by as seen from the target. In the (u, v) plane, the array is projected to two dimensions, which results in baselines being shortened along the v direction by a factor of $\cos(\theta_{el})$. Figure 5.8 shows the (u, v) -coverage of a typical MALS observation of a single target at several different declinations, showing the projection effect. As MeerKAT has a latitude of -30° , we expect that the array is least affected by this projection if the source is at the same declination.

5.2.1 (u, v) -coverage effects

The projection of the array and consequent change in (u, v) -coverage for observations at different declinations carries through to the shape of the synthesised beam of the image. As the baselines are shortened in the projected array in one direction, the spatial resolution in that direction is decreased. This can be quantified by the baseline length in the v direction, as its shown in Figure 5.8 that this is the affected quantity at different declinations. More direct measures of (u, v) -coverage are not commonly used to quantify radio observations, as it depends on the details of imaging. We here additionally define the (u, v) -coverage factor, f_{uv} , as the fraction of pixels in the (u, v) -grid that contain at least one measured visibility. The pixel size of the (u, v) -grid is determined by the size of the image, as

$$\delta u = \delta v = \frac{1}{\theta_{fov}}. \quad (5.3)$$




Figure 5.7: Number of sources per pointing, plotted for each pointing separately as a function of their declination (top) and averaged over different declination bins (bottom), for different cuts in flux density. Additionally, the mean number of sources per pointing is shown by the dotted line for each flux density cut. The edges of the declination bins are indicated by the solid grey lines. The lowest cut includes all sources, and so goes down to around $50 \mu\text{Jy}$.

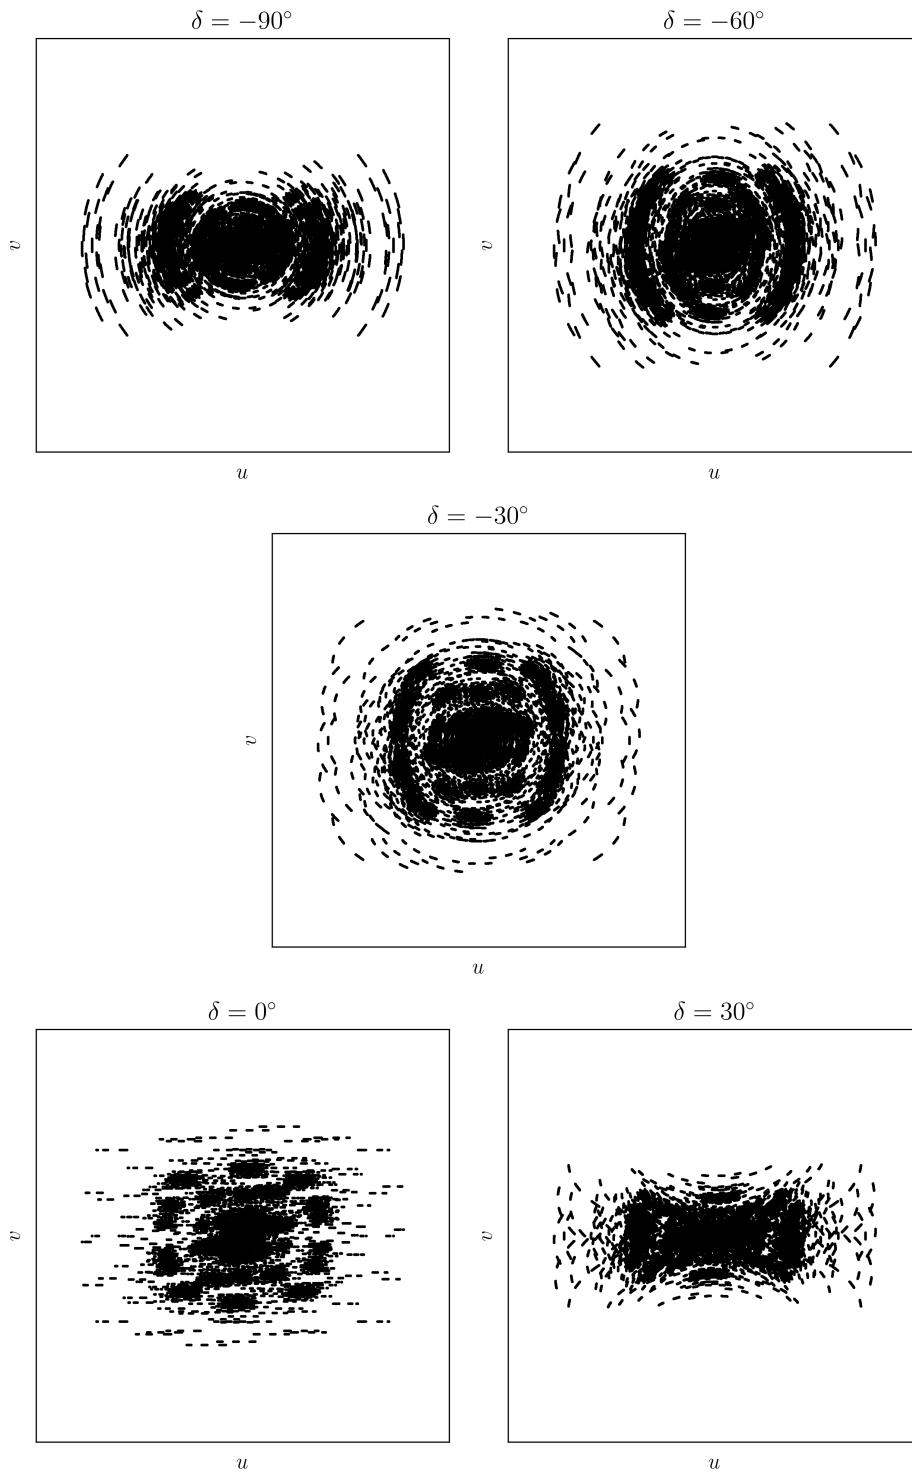


Figure 5.8: Example (u, v) -coverages of MALS-like observations at different declinations. The most circular (u, v) -coverage exists at declinations closer to the latitude of MeerKAT, where the elevation of the target would be higher.




Figure 5.9: Quantities from (u, v) -coverages of simulated MALS pointings as a function of declination. Both maximum baseline length in the v direction (left) as well as (u, v) -coverage factor f_{uv} (right) are shown, along with the best fit absolute cosines (solid red line, Equation 5.4). The relation fits well in both cases, but slightly deviates at higher declinations for f_{uv} .

In the case of our MALS images, with a field of view of 3.3° , the corresponding pixel size in the (u, v) -grid is 18λ .

In order to determine both f_{uv} and the maximum baseline in the v -direction for the MALS pointings, we simulate (u, v) -coverages, assuming that all pointings are observed three times in the span of three hours, 20 minutes at a time, emulating the original observing setup. We furthermore assume that the central target reaches zenith in the middle of the observations. This produces (u, v) -coverages such as the ones shown in Figure 5.8. The associated properties of the (u, v) -coverages are shown in Figure 5.9. The left plot of Figure 5.9 shows the variation of the maximum baseline in the v -direction as a function of declination, which resembles an cosine as a function of declination. The right plot of Figure 5.9 shows the simulated (u, v) -coverage factors of all MALS pointings, showing a very clear evolution of f_{uv} as a function of declination. It however peaks at a different declination than that corresponding to the latitude of the telescope. This declination dependence strongly resembles declination dependences seen in some other surveys, which we will discuss in Section 5.6.

The relation to declination seen in both these quantities resemble a function of the form

$$f(x) = a_0 + a |\cos[c(x - b)]|. \quad (5.4)$$

Fitting this function to the simulated maximum v baseline as a function of declination, we find a minimum baseline $a_0 = 10.6 \text{ k}\lambda$, with amplitude $a = 8.7$, offset $b = -30$, and width $c = 1.53$. Fitting this function to the simulated f_{uv} values as a function of declination, we find a minimum coverage $a_0 = 0.0074$, with amplitude $a = 0.012$, offset $b = -83$, and width $c = 1.08$. Figure 5.9 shows that these fits agree well with the simulated data, bit slightly deviates for f_{uv} at the highest declinations. As such, we have two quantities that could describe the declination dependency of the source counts. While both have the same cosine shape, they peak at different declinations. This uncertainty in where the peak lies must be taken into account if we wish to model the effect in our dipole estimates.




Figure 5.10: Number counts of all MALS pointings (black points) as a function of declination, along with the best fit absolute cosine (red line, Equation 5.5).

5.2.2 Fitting the data

Though we aim to fit the declination effect as part of the dipole estimation, we can get an initial estimate of the effect by fitting Equation 5.4 to the data. To fit this function to number counts, we define

$$N(\delta) = N_0\{1 + a|\cos[c(\delta - b)]|\}, \quad (5.5)$$

with N_0 the base number counts, and all other parameters as before in Equation 5.4. With this, we fit the effect to the number of sources per pointing as a function of declination without any cuts in flux density. Figure 5.10 shows the data and resulting fit, with the best fit cosine having an amplitude $a = 1.38$, offset $b = -52$, and width $b = 1.54$. The resulting offset does not favour solely the maximum ν -baseline or f_{uv} parameter to describe the declination dependence, implying that it might be a combination of several effects. This again emphasises the importance of leaving the offset, b , a free parameter. It is however important to note that this fit favours a minimum at $\delta = 0$, which is more in line with a f_{uv} effect.

5.3 Number counts

With a catalogue containing nearly 8×10^5 sources, we can obtain robust differential number counts covering nearly five orders of magnitude of flux density. For number counts of bright sources, there are a number of shallow surveys that cover a large area of the sky, such as NVSS (Condon et al., 1998) and RACS (Hale et al., 2021). Number counts of faint sources are usually obtained with deep observations of small areas such as DEEP2, COSMOS, or XMM-LSS (e.g. Smolčić et al., 2017; Mauch et al., 2020;




Figure 5.11: Differential number counts of MALS, uncorrected (beige circles) and corrected with rms coverage (green squares). Counts are compared to SKADS (Wilman et al., 2008) and the combined source counts from NVSS and DEEP2 (Matthews et al., 2021b).

van der Vlugt et al., 2021; Hale et al., 2023). These number counts can be combined to cover a large range in flux density as is done in Matthews et al. (2021b), however these measurements might not perfectly connect and completely cover the full range of number counts. As MALS strikes a balance between depth and sky coverage, our differential number counts should cover the transition between large sky shallow surveys and small sky deep surveys.

In Chapter 3, we produced differential number counts of the first ten MALS pointings (Figure 3.24). For these ten pointings, we showed that though 100% completeness was only reached around 1 mJy, we could apply corrections to the number counts based on completeness measures to be complete down to 200-300 μ Jy. We also noted that differential number counts were higher than expected, however with just ten pointings still some uncertainty remained (Wagenveld et al., 2023b). Figure 5.11 shows the differential source counts for all 391 MALS pointings, giving a significant increase in precision compared to only the first ten pointings. The counts remain consistent however, and with the rms coverage corrections the number counts are complete down to 300 μ Jy. Furthermore, it is now readily apparent that the MALS number counts are increased with respect to both the theoretical number counts from SKADS (Wilman et al., 2008) and the observed number counts from NVSS and DEEP2 (Matthews et al., 2021b). This difference persists for the entire range of flux densities, and is even further increased at flux densities above 200 mJy. This last effect is likely due to the pointing selection of MALS, which always contains a source with flux density above 200 mJy in the pointing centre. The increased number counts in the entire range of flux densities could be caused by the same selection effect selecting for overdensities. On the other hand, having bright central sources has




Figure 5.12: Differential number counts of MALS, uncorrected (circles) and corrected with rms coverage (squares) for declinations below (blue) and above (red) $\delta = -20^\circ$. Counts are compared to SKADS (Wilman et al., 2008) and the combined source counts from NVSS and DEEP2 (Matthews et al., 2021b).

allowed for (improved) self-calibration of the MALS data, so perhaps the number counts we find are more accurate than the ones from surveys which lack this advantage. Though the effect is noteworthy, it merely increases the monopole and is not expected to influence a dipole measurement.

At the lowest flux densities, there is a well established disagreement between the simulated sample of star-forming galaxies from SKADS and observed samples. In observations, the amount of star-forming galaxies needed to explain the observed number counts is generally much larger than the amount predicted in SKADS (Matthews et al., 2021b; Hale et al., 2023, e.g.). In Figure 5.11 we see that at these flux densities the MALS number counts are more consistent with SKADS in terms of their evolution with respect to flux density, however the increased source density with respect to SKADS is more in agreement with the Matthews et al. (2021b) number counts.

To investigate if the observed declination variation has any influence on the number counts, we plot the number counts in two different declination bins in Figure 5.12. Though number counts are only significantly affected for the fainter source population the effect, however subtle, does persist above the completeness limit of the survey. Although the rms coverage corrections mitigate the effect somewhat, it is still significant at the flux densities where we consider the corrected number counts complete. Though source counts are increased for data at $\delta < -20^\circ$, especially at low flux densities, they still follow SKADS number counts more closely than the Matthews et al. (2021b) number counts. The fact that the declination effect is flux density dependent would mean that to properly correct the number counts for the effect, it would have to be fit for each flux density bin. As such, the differential number counts in Figure 5.11 are not corrected for the effect, adding an additional source of uncertainty at the lowest flux densities. This effect however is arguably only significant—for the differential number counts specifically—below the (corrected) completeness limit of (300 μ Jy) 1 mJy.

5.4 Preparing for a dipole measurement

Unlike most surveys used for measurement of the number count dipole, MALS sparsely populates the sky with deep pointings, rather than contiguously covering the sky with more shallow observations. This means that in order to measure the dipole with MALS, dipole estimators are required which are unbiased by gaps in the data. Advantageously, MALS does not require pixelisation, as every pointing covers the same amount of sky area. As such, number counts can be obtained simply by counting the number of sources in each pointing.

5.4.1 MALS data

As mentioned previously, a number of pointings have high noise and therefore low number counts, as shown in Figure 5.6. For a dipole measurement, we discard all of these pointings. Though there are 41 high noise pointings, only a small fraction of all sources ($\sim 5\%$) are removed. To avoid counting Galactic sources, we remove the five pointings at low galactic latitudes ($|b| < 10^\circ$). This leaves 345 pointings for a dipole estimate, with a total of $\sim 7.5 \times 10^5$ sources.

Figure 5.11 shows that the 100% completeness limit required for a homogeneous catalogue is only reached around 1 mJy. Less than 2×10^5 sources are present in the full catalogue above 1 mJy, which is insufficient for a dipole measurement. However, using completeness corrections, the catalogue can be made complete down to much lower flux densities. Figures 5.11 and 5.12 show that these corrections might even hold down to around 150 μJy , however to be slightly conservative we restrict ourselves to a minimum flux density of 300 μJy . Down to this level, the catalogue contains $\sim 4 \times 10^5$ sources, which is sufficient for a significant measurement of the dipole, assuming a dipole amplitude similar to that found in previous measurements of the radio dipole (e.g. Siewert et al., 2021; Secrest et al., 2022; Wagenveld et al., 2023b, and Chapter 4).

We compute the completeness of sources in two different ways. The first method utilises the rms coverage curves shown in Figure 5.3. For each source, we determine the associated rms based on the peak flux density of the source divided by five (matching the 5σ detection threshold). The completeness is then assigned based on the coverage of the associated rms value in that pointing. The second method uses the completeness obtained from the simulations performed in Chapter 3. The combined completeness functions are determined separately for resolved and unresolved sources, as a function of σ_{20} of the pointing, and are shown in Figure 3.15. These corrections essentially measure for all sources in what fraction of the sky coverage of the pointing they would be detected.

Following Equation 1.1 we determine the expected dipole amplitude of MALS, assuming a velocity equal to that measured from the CMB dipole ($v = 370$ km/s, Aghanim et al., 2020). In Chapter 3 we investigated the spectral index distribution of MALS sources, showing that down to 1 mJy the mean spectral index is $\langle \alpha \rangle = 0.76$. We measure the power law index of the flux density distribution of MALS to be $x = 0.68$. Though the power law index could vary as a function of flux density as the underlying source population at the measured flux density (300 μJy) consists of both AGN and SFGs, this power law index is valid at flux densities between $0.3 \leq S \leq 100$ mJy, as is shown in Figure 5.13. For an amplitude equal to that of the CMB, this sets the expectation of a dipole amplitude measurement to $\mathcal{D} = 0.39 \times 10^{-2}$.




Figure 5.13: Flux density distribution of MALS sources above 300 μJy , along with fitted power law. The power law fits the flux density distribution well up to 100 mJy.

5.4.2 Estimators

In Chapter 4, we defined a set of Bayesian estimators based on Poisson statistics that should be insensitive to incomplete sky coverage and will thus work with the sparse structure of MALS. These estimators assume that the number of sources in the pointings will follow a Poisson distribution. The associated likelihood is

$$\mathcal{L}(n) = \prod_i \frac{\lambda^{n_i} e^{-\lambda}}{n_i!}, \quad (5.6)$$

where n_i is the amount of sources in pointing i and λ describes the mean and variance of the distribution. For the basic Poisson estimator

$$\lambda(\vec{d}, \mathcal{M}) = \mathcal{M}(1 + \vec{d} \cdot \hat{n}), \quad (5.7)$$

where \mathcal{M} is the monopole or the mean amount of source per pointing. The dipole is described by the dipole vector \vec{d} . The effect of the dipole is based on location on the sky, $\vec{d} \cdot \hat{n} = \mathcal{D} \cos(\theta)$ with θ the angular distance between the dipole direction and the direction of the pointing.

The Poisson-rms estimator aims to model the relation between the local noise to the source density, which removes the necessity for a cut in flux density. Figure 5.6 shows the number of sources in each pointing as a function of σ_{20} , along with a power law fit, showing this model fits the MALS data as well. As such, utilising the Poisson-rms estimator, we can use the full catalogue of sources for a dipole measurement without needing to apply a cut in flux density. The relation between source density and

rms that is assumed by the Poisson-rms estimator is given by

$$\lambda(\vec{d}, \mathcal{M}, \sigma, x) = \mathcal{M} \left(\frac{\sigma}{\sigma_0} \right)^{-x} (1 + \vec{d} \cdot \hat{n}). \quad (5.8)$$

Here, σ is the noise associated with the pointing, x is the power law index, and σ_0 is the normalisation factor for the noise. We determine σ_0 as the median σ_{20} of the included pointings, which sets it to 25.8 $\mu\text{Jy}/\text{beam}$.

Finally, we define the declination estimator, an additional estimator which aims to take into account the systematic variation in source density as described in Section 5.2. For most plausible explanations of the variation in source density, this effect is expected to follow a cosine function as a function of declination, δ . Combined with the dipole, the source density is affected

$$\lambda(\vec{d}, \mathcal{M}, \delta, a, b) = \mathcal{M} [1 + a |\cos(\delta - b)|] (1 + \vec{d} \cdot \hat{n}), \quad (5.9)$$

where a is the amplitude of the variation, and b is the offset. The width parameter c that was present in Equation 5.5 is dropped here, in favour of simplifying the function and decreasing the amount of free parameters.

The Bayesian estimators described here have been implemented using the Bayesian inference library BILBY (Ashton et al., 2019). Through BILBY, we maximise the likelihood with MCMC sampling using EMCEE (Foreman-Mackey et al., 2013). The chosen priors and injection values are the same as those described in Chapter 4.3.5. For the estimator including the declination dependence, we constrain the amplitude and offset of the cosine with uniform priors $\pi(a) = \mathcal{U}(0, 1)$ and $\pi(b) = \mathcal{U}(-90, 0)$. The injection values for these parameters are $a = 0.01$ and $b = -45$.

5.4.3 Evaluating the estimators with simulations

To examine how our estimators perform on a catalogue such as MALS, we simulate a set of catalogues with the same sky distribution of MALS. To generate source positions we uniformly distribute sources within a radius of 1.15° of all MALS pointing centres. We generate rest flux densities S_{rest} according to the power law

$$S_{rest} = S_{low} (1 - \mathcal{U})^{-1/x}, \quad (5.10)$$

where \mathcal{U} is a uniform distribution between 0 and 1, S_{low} is the lower flux density limit at which sources are generated, and x the power-law index of the flux density distribution. All simulated sources have an assigned spectral index of 0.75, and the power law index of the flux density distribution is set to $x = 1$.

On this generated source population, we simulate a dipole effect. To transform rest flux densities to the frame of the moving observer, we apply a Doppler shift and a Doppler boost as described in Equation 4.4. We apply relativistic aberration to transform the source positions to the frame of the moving observer as expressed in Equation 4.5. We set the direction of movement close to the direction of the CMB dipole, (R.A., Dec.) = $(170^\circ, -10^\circ)$. We apply these effects with an increased velocity of $v = 1200 \text{ km/s}$ ($\beta = 4 \times 10^{-3}$) in order to more closely match previously measured amplitudes of the radio dipole, and to require less sources for a significant measurement. Following Equation 1.1, this sets the expected dipole amplitude to $\mathcal{D} = 1.5 \times 10^{-2}$.

To create the observed catalogues, we apply Gaussian noise to the flux densities and only retain

Table 5.1: Dipole estimates using the various estimators on the simulated catalogues.

Simulation	Estimator	Correction	S_0 (μJy)	N	\mathcal{M} counts/pointing	x	\mathcal{D} ($\times 10^{-2}$)	R.A. (deg)	Dec. (deg)
i	Poisson	–	300	1,199,394	3427 ± 4	–	0.71 ± 0.14	171 ± 13	6 ± 22
		–	500	717,063	2049 ± 3	–	0.56 ± 0.21	177 ± 25	19^{+29}_{-36}
ii	Poisson-rms	–	–	1,496,022	4152 ± 5	0.999 ± 0.004	0.80 ± 0.14	154 ± 10	-6 ± 21
	Poisson	–	300	602,177	1722 ± 3	–	0.88 ± 0.24	147^{+18}_{-15}	28^{+20}_{-25}
	–	–	500	358,822	1025 ± 2	–	0.94 ± 0.30	140^{+20}_{-17}	-3^{+31}_{-29}
iii	Poisson-rms	–	–	3,683,920	10226 ± 7	1.002 ± 0.002	0.74 ± 0.08	166 ± 8	0 ± 17
	Poisson	–	300	2,034,506	5615 ± 5	–	14.5 ± 0.2	151 ± 12	-88 ± 1
		–	500	1,372,769	3851 ± 4	–	7.6 ± 0.3	168 ± 12	-85 ± 1
		–	1000	718,765	2051 ± 3	–	0.78 ± 0.22	171 ± 15	-28^{+22}_{-17}
	Poisson	SIM	300	2,034,506	5615 ± 5	–	2.3 ± 0.19	172 ± 9	-72 ± 3
		SIM	500	1,372,769	3851 ± 4	–	$2.9^{0.24}_{-0.35}$	170^{+15}_{-11}	-73^{+4}_{-3}
	Poisson	RMS	300	2,034,506	5615 ± 5	–	1.06 ± 0.14	184 ± 8	44^{+7}_{-9}
		RMS	500	1,372,769	3851 ± 4	–	1.01 ± 0.14	173 ± 9	27^{+12}_{-15}

sources with $S/N > 5$. In order to disentangle potential effects affecting the dipole measurement, we create three simulated catalogues with different noise properties. The noise properties of these catalogues are (i) the same rms ($\sigma = 20 \mu\text{Jy}/\text{beam}$) for all pointings, (ii) the measured σ_{20} noise level throughout each pointing, (iii) noise to each source according to its position within the pointing, using the median rms map shown in Figure 5.4, scaled to the σ_{20} value of that pointing.

Results

The results of the dipole estimates on the simulated catalogues are summarised in Table 5.1. Due to the different noise structures in the simulated catalogues, not all estimators and corrections can be applied to all catalogues. In summary, the estimates that are affected by the incompleteness of the survey are highly biased towards the south pole, where the noise is lower due to the declination effect described in Section 5.2. These estimates also have much higher amplitude, indicating the strength of the effect. For the estimates where the the correct direction is recovered within 3σ , the amplitude is systematically underestimated, usually not exceeding $\mathcal{D} = 1.0 \times 10^{-2}$ where the expected value is $\mathcal{D} = 1.5 \times 10^{-2}$.

The most realistic representation of the noise structure in the survey is simulation (iii), where we use the median rms map from Figure 5.4 to generate the local noise. Table 5.1 shows that the Poisson-rms estimator correctly estimates the dipole direction, but underestimates the dipole amplitude by a factor of two. For different flux density cuts, we see that incompleteness causes estimates with flux density cuts below 1 mJy to be biased towards the south pole. For this simulated catalogue, we can also apply the completeness corrections described in Section 5.4. We see that the completeness corrections from the simulations (SIM) do not completely alleviate the bias introduced by incompleteness of the catalogue, and the dipole estimates are still biased towards the south pole. The completeness corrections from the rms coverage (RMS) actually seem to overcompensate for the incompleteness,




Figure 5.14: Distribution of estimated dipole amplitudes for the different estimator on 100 simulated catalogues. The counts used for the estimates with 300 μJy and 500 μJy flux density cuts are corrected with the rms coverage. The expected dipole amplitude is indicated by the solid black line.

slightly biasing the dipole estimate to the north. The dipole amplitudes for these estimates are however much closer to the true value. We note that in all cases, the right ascension estimates match the input value within 3σ , showing that there is no inherent bias due to incompleteness in right ascension. These simulations show that the Poisson-rms estimator should most reliably yield the correct (though biased) result. A flux density cut of 1 mJy would also yield a good estimate, however as mentioned before, less than 2×10^5 sources are present above this flux density in the MALS catalogue. For most of these estimates, the dipole amplitude is underestimated by a factor of two on average.

Quantifying the amplitude bias

In order to quantify the bias seen in the dipole amplitude estimates, we perform simulation (iii) 100 more times perform dipole estimates. Each catalogue contains around 4×10^6 sources, such that we can make cuts above the completeness limit while still retaining enough sources for a dipole estimate. On each catalogue we estimate the dipole with the Poisson-rms estimator, and with the regular Poisson estimator for different flux density cuts. The counts are corrected for the rms coverage for the 300 μJy and 500 μJy flux density cuts, while no corrections are made for the 1 mJy flux density cut.

Figure 5.14 shows the distribution of estimated dipole amplitudes for the different estimators and flux density cuts. We see that in all cases the dipole amplitude is underestimated. The amplitudes recovered by the 1 mJy flux density cut and Poisson-rms estimator are consistent, with mean amplitudes of $\mathcal{D} = (0.74 \pm 0.08) \times 10^{-2}$ and $\mathcal{D} = (0.78 \pm 0.16) \times 10^{-2}$ respectively. As on average five times fewer sources are available after the 1 mJy flux density cut, the variance in amplitudes is higher for these estimates. In both cases, the amplitude is consistently underestimated by a factor of two on average. As shown in Table 5.1, the estimates using number counts corrected for the rms coverage are

5.4 Preparing for a dipole measurement

Table 5.2: Dipole estimates using the various estimators and flux density thresholds for different subsets of the MALS data. Estimates that did not converge on a solution have been omitted.

Subset	Estimator	Correction	S_0 (μJy)	N	\mathcal{M} counts/pointing	x	\mathcal{D} ($\times 10^{-2}$)	R.A. (deg)	Dec. (deg)
All	Poisson-rms	–	–	724,566	1948 ± 3	1.085 ± 0.006	13.89 ± 0.50	125 ± 4	-79 ± 1
	Poisson	RMS	300	398,284	1710 ± 3	–	2.34 ± 0.30	182 ± 8	-43 ± 8
		RMS	500	277,450	959 ± 2	–	3.51 ± 0.51	134 ± 14	-71 ± 5
		–	1000	166,572	476 ± 1	–	6.17 ± 0.75	148_{-19}^{+23}	-79 ± 4
$\delta < 10$	Poisson-rms	–	–	657,728	2039 ± 5	1.118 ± 0.006	7.51 ± 0.60	119 ± 5	-73 ± 2
	Poisson	RMS	300	354,689	1737 ± 4	–	2.03 ± 0.43	196 ± 17	66_{-9}^{+7}
		RMS	500	246,110	974 ± 2	–	$0.57_{-0.36}^{+0.42}$	126_{-36}^{+55}	22_{-59}^{+45}
		–	1000	147,142	483 ± 2	–	2.73 ± 0.98	136_{-30}^{+41}	-74_{-9}^{+12}
$\delta < -10$	Poisson-rms	–	–	474,531	2432 ± 9	1.145 ± 0.007	2.78 ± 0.35	112 ± 5	-17 ± 14
	Poisson	RMS	300	235,481	1788 ± 8	–	$7.17_{-0.82}^{+0.74}$	246_{-29}^{+25}	85 ± 3
$\rho < 0.5^\circ$	Poisson-rms	–	–	301,842	805 ± 2	1.265 ± 0.009	11.57 ± 0.72	128 ± 7	-78 ± 2
	Poisson	RMS	300	94,693	552 ± 2	–	30.03 ± 0.58	223 ± 7	84 ± 1
		–	1000	29,409	85 ± 1	–	$3.95_{-1.42}^{+1.63}$	138_{-19}^{+24}	-46_{-17}^{+24}
Unresolved	Poisson-rms	–	–	391,917	1039 ± 2	1.353 ± 0.008	11.44 ± 0.64	136 ± 4	-68 ± 2
	Poisson	RMS	300	138,258	566 ± 2	–	9.34 ± 0.61	170 ± 5	60 ± 3
		RMS	500	76,379	257 ± 2	–	12.04 ± 0.84	183 ± 5	58 ± 3
		–	1000	32,822	99 ± 1	–	17.12 ± 1.37	195 ± 7	63 ± 4
No clusters	Poisson-rms	–	–	653,492	1957 ± 4	1.087 ± 0.007	13.17 ± 0.58	133 ± 5	-78 ± 1
	Poisson	RMS	300	357,030	1710 ± 3	–	1.82 ± 0.26	172 ± 8	-27 ± 11
		RMS	500	248,457	962 ± 2	–	2.27 ± 0.51	127 ± 12	-59 ± 8
		–	1000	148,916	477 ± 2	–	5.16 ± 0.79	139 ± 22	-78 ± 5

biased north and yield higher dipole amplitudes with respect to complete estimates. Similar results are seen here, with the 300 μJy and 500 μJy flux density cuts yielding mean dipole amplitudes of $\mathcal{D} = (1.06 \pm 0.15) \times 10^{-2}$ and $\mathcal{D} = (1.00 \pm 0.19) \times 10^{-2}$ respectively. Thus, these underestimate the dipole amplitude by a factor of 3/2 on average.

The cause for these biases is not quite clear, as the estimators should be insensitive to gaps in the data. However, as it persists for different simulated noise structures, it is likely a result of the sky coverage of MALS with respect to the chosen dipole axis, for example by improper sampling of the dipole axis by the MALS pointings. If this is the origin of the bias, the bias might be different if the pointings were distributed differently, or if there were more or less pointings. To determine this would require many more simulations, changing the number and location of the pointings. This is beyond the scope of the current work, as for now determining the bias for the specific sky coverage of MALS is sufficient. This does however show that simulations provide an important check for unexpected biases in any estimator.

Table 5.3: Dipole estimates on the MALS catalogue with the declination estimator described in Equation 5.9, for various flux density cuts and completeness corrections. In addition to the dipole parameters, parameters of the cosine (amplitude a and offset b) describing the declination dependent source density are also estimated. As specified, a Gaussian is used as prior distribution on declination of the dipole direction (see also Figure 5.16).

Correction	S_0 (μJy)	N	\mathcal{M} counts/pointing	a	b (deg)	\mathcal{D} ($\times 10^{-2}$)	R.A. (deg)	Dec. (deg)
RMS	300	398,284	1497 ± 24	0.16 ± 0.02	-24 ± 3	$1.57^{+0.25}_{-0.21}$	183 ± 9	0 ± 21
RMS	500	277,450	810 ± 18	0.21 ± 0.03	-38 ± 2	$0.97^{+0.37}_{-0.35}$	124 ± 17	2 ± 22
SIM	300	398,284	1220 ± 23	0.29 ± 0.02	-46 ± 2	$1.28^{+0.29}_{-0.26}$	139 ± 11	1 ± 22
SIM	500	277,450	778 ± 18	0.24 ± 0.03	-42 ± 3	$1.48^{+0.38}_{-0.34}$	120 ± 10	0 ± 21
–	1000	166,572	417 ± 13	0.19 ± 0.04	-39 ± 5	$1.01^{+0.46}_{-0.43}$	142 ± 24	-1 ± 21

5.5 MALS results

With the cuts in the data specified in Section 5.4, removing high noise pointings and pointings with low galactic latitude, we now perform dipole estimates on the MALS catalogue. A selection of results with different estimates and additional cuts in the data is given in Table 5.2. From Table 5.2 it is clear that all estimates are heavily affected by the systematic variation in source density as a function of declination. This effect dominates the dipole estimation for MALS on both the Poisson and Poisson-rms estimators. Different flux density cuts, as well as different declination cuts, and even cuts in terms of distance from the pointing centre do not alleviate this effect. Though the Poisson-rms estimator was not affected by the noise structure in the simulated data sets, on the MALS data it, along with the other estimators, consistently yields biased results.

5.5.1 Declination estimator

As the basic Poisson and Poisson-rms estimators all yield results in which the declination is heavily biased towards the south pole, we employ the declination estimator which takes a declination effect into account. Figure 5.15 shows the posterior distributions of all fitted parameters for such an estimate, using a flux density threshold of 300 μJy . This shows that the estimator is able to fit the declination relation present in the data, along with the dipole parameters, without any additional cuts required. However, as the declination dependence and dipole have the same functional shape, the peak of the declination function b and the declination of the dipole are highly degenerate. As a consequence, the declination of the dipole becomes mostly unconstrained, being nearly equally likely for all declinations in the range $(-60,60)$. As is seen in Figure 5.15, other parameters are more poorly constrained as well as a result of this. We do see however that the maximum likelihood values of the dipole parameters closely resemble previously measured values of the radio dipole, showing that this method can account for the declination effect.

To better constrain the fit parameters, we must put tighter constraints on the range of possible declinations of the dipole. In order to do this, we change the prior on the declination of the dipole to reflect our disbelief of the dipole being present at higher declinations. Instead of a uniform prior on




Figure 5.15: Posterior distributions for all parameters fit with the declination estimator (Equation 5.9, dipole amplitude \mathcal{D} , monopole \mathcal{M} , declination cosine amplitude a and offset b), for a flux density threshold of $300 \mu\text{Jy}$. The contours indicate 1-,2-, and 3- σ uncertainties. Due to the declination dependence being fit with a cosine, the declination of the dipole has remained completely unconstrained. This impacts the estimates of other parameters as well, as for example at high latitudes higher dipole amplitudes are preferred. The best fit parameters for the combined RACS and NVSS paper from Chapter 4 are also shown (green lines), as well as the direction of the CMB dipole (orange lines).

the whole range of possible dipole declinations, we adopt a Gaussian prior centred on the equator; $\pi(\text{Dec.}) = \mathcal{G}(\mu = 0, \sigma = 20)$. Figure 5.16 shows the posterior distributions of all fitted parameters with this prior, using the same flux density threshold of $300 \mu\text{Jy}$. We see that as a result of choosing a more constrained prior, other parameters are more constrained as well. Table 5.3 shows the parameter estimates for multiple flux density cuts and corrections using this prior and the declination estimator.

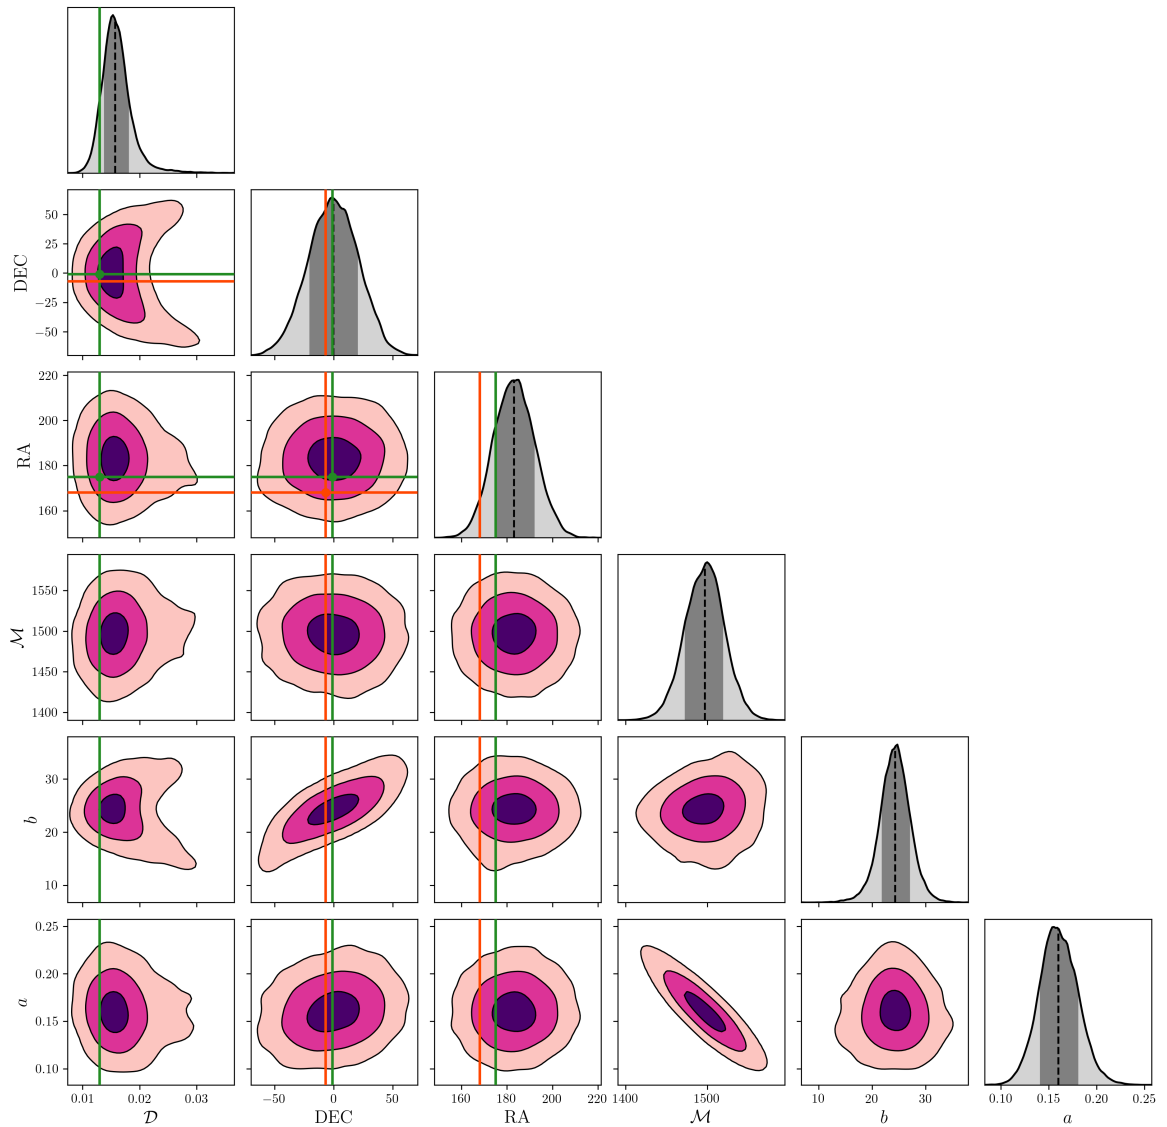


Figure 5.16: Posterior distributions for all parameters fit with the declination estimator, for a flux density threshold of $300 \mu\text{Jy}$. The contours indicate 1-,2-, and 3- σ uncertainties. With a Gaussian prior on the declination of the dipole direction, the constraints have improved for all other parameters. The best fit parameters for the combined RACS and NVSS paper from Chapter 4 are again shown (green lines), as well as the direction of the CMB dipole (orange lines).

In terms of the declination dependence, for most estimates the function peaks at a declination of $\delta \approx -40^\circ$, and decreases in amplitude at higher flux density cuts, indicating that sources at higher flux density are less affected by the declination systematic. An exception is the $300 \mu\text{Jy}$ flux density cut with rms correction, for which the function peaks at a declination of $\delta = -24^\circ$, and also has a lower amplitude.

We see that in all cases estimated dipole parameters are broadly consistent with previous results

of the radio dipole. The most precise estimates are obtained with a flux density cut of $300 \mu\text{Jy}$, also shown in Figure 5.16, with the estimate using the rms correction rejecting the CMB dipole amplitude (as determined in Section 5.4) with a formal significance of 5.6σ . The fact that this estimate also is the closest in terms of direction to the CMB dipole lends it extra credibility. The other estimate with $300 \mu\text{Jy}$, using the simulation correction, is closer to previous dipole estimates with NVSS (e.g. Siewert et al., 2021; Wagenveld et al., 2023b), both in terms of direction and amplitude, with a significance of 3.4σ . Other estimates have too few sources for a significant detection, but we note that amplitudes and directions remain broadly consistent. How these estimates have now been affected by the amplitude bias discussed in Section 5.4 is not very clear, as the addition of the declination dependence to the parameter estimation could have affected this bias in unexpected ways. This leaves some ambiguity in the obtained dipole amplitudes, however we can be certain that the dipole amplitudes are no smaller than the ones obtained here, which are all larger than the dipole amplitude expected from the CMB by at least a factor of two.

5.6 Discussion

The declination effect seen in the MALS data dominates the dipole signal and introduces an artificial dipole effect which points towards the south pole with high amplitude. Though a variation of source density in declination was already seen in the noise distribution of the pointings, the Poisson-rms estimator and completeness corrections, which should explicitly account for the noise, failed to alleviate the bias. It is clear then, that the effect in question goes deeper than the noise. In Section 5.2 we explored possible causes of the declination dependence, with the most likely scenario being related to the projection of the array during observations. With this we showed that different effects related to the (u, v) -coverage all follow a similar shape, but the declination at which they peak might be different. The results in Table 5.3 indicate that the cosine describing the declination dependence, with a single exception, peaks around -40° . An effect purely following the beam shape would be expected to peak at -30° , while an effect purely following the (u, v) -coverage factor would peak at around -80° . This might indicate that the true cause might be a combination of these, one factor perhaps affecting source density more than the other.

Though we have shown that these systematics can be taken into account during parameter estimation, the results show that these effects, which have the same shape as a dipole, prevent a well constrained measurement of the declination of the dipole. It is difficult to imagine a model in which the declination dependence is not degenerate with the declination of the dipole. Therefore, the ultimate way to alleviate this systematic is to reprocess the data to remove the effect causing the systematic in the first place. If the effect causing the source variation is related to the (u, v) -coverage, either through the shape of the beam or the coverage fraction of the (u, v) grid, a few solutions present itself to mitigate the effect.

One solution is used by large contiguous radio surveys, which is to make all images have a common resolution. After imaging this can be achieved by smoothing of the existing images, or during imaging by tapering in the (u, v) plane. Naturally, this common beam has to be equal or larger than the largest beam of an individual pointing. As by doing this we accept the premise that a larger beam decreases the amount of detected sources in an image, the homogeneity comes at a cost of less sources in the overall catalogue. One possibility to mitigate this loss in source density over the entire catalogue is to define several declination bins in which the data is smoothed by the largest beam in that bin. A

dipole measurement could then be performed utilising the multiple catalogue estimator introduced in Chapter 4.

5.6.1 Declination systematics in other radio surveys

The main crippling factor for a dipole estimate with MALS was its relatively high depth and small sky coverage, as a systematic affecting mostly fainter sources is greatly enhanced. Though this is the one of the first works to actively tackle such an overwhelming systematic, MALS is certainly not the only survey suffering such an effect, and similar effects can be seen in other radio surveys. A similar to the one we see in MALS appears in LoTSS (D. Schwarz, private communication, 2023). Though in Chapter 4 we were able to alleviate the RACS systematic for our Poisson estimator with a high enough flux density cut, the source density of the faintest sources strongly varies as a function of declination. This affected the result of the Poisson-rms estimator, which shows that as with the systematic effect seen here, the effect likely goes beyond variations in local noise. As shown in Figure 4.1, the source density peaks at $\delta \approx -75$ and troughs at $\delta \approx 0$, showing a variation remarkably similar to the relation between declination and f_{uv} shown in the right plot of Figure 5.9. These effects could have affected the strange result of Darling (2022), who combined VLASS and RACS for a dipole measurement and found a dipole that was actually consistent with CMB, however uncertainties were so large that the measurement was also consistent with other measurements which reject the CMB dipole. A lack of a measurement on the individual catalogues makes this however difficult to assess. Using a linear estimator, Singal (2023) did measure the dipole with VLASS and RACS separately, and both estimates were adversely affected by the aforementioned systematics in these catalogues.

These instances are worth mentioning, as these effects are still present in these catalogues despite them being imaged and smoothed properly. In fact, some of these effects similarly persist far above the completeness limit of these catalogues. We can conclude that even with great care taken in data processing and analysis, such systematics can be extremely persistent, and the best approach is to mitigate any systematics as much as to allow a unbiased dipole estimate. Whether this can be achieved in full with MALS with the additional suggested processing steps, remains to be seen.

5.7 Conclusion

In this work we have presented a first attempt to measure the cosmic radio dipole on MALS. For this we have compiled one of the largest radio catalogues currently available, comparable in size to NVSS, RACS, and VLASS. Having 391 observed pointings in L band, the full MALS survey has enough sky coverage and depth for a catalogue containing hundreds of thousands of sources, sufficient for a significant measurement of the radio dipole. We construct a catalogue for a dipole measurement by including all sources within a radius of 1.15° of the pointings centres, and flagging sources that are likely to be artefacts. This leaves us with a survey coverage of 1623 deg^2 and a catalogue containing 796,916 sources. We show that all pointings have similar noise structures, and that we can describe their differences with σ_{20} , the rms value at a cumulative area coverage of 20%. We see that there is a tight power law relation between σ_{20} and the number of sources in a pointing. Using corrections defined in Chapter 3, we use several completeness corrections for the catalogue at lower flux densities to increase the number of sources that can be used for a dipole estimate.

In the data we see a systematic variation in source density which varies as a function of declination.

This effect goes beyond just a noise variation, and persists above the completeness limit of the survey. As the observed systematic induces the variation in source density only in declination, it likely originates from effects of projection of the array. We characterise the change in two quantities measured from the change in (u, v) -coverage of the array as a function of declination of the source, finding a common function to describe these quantities. We introduce an additional dipole estimator, which aims to account for the variation in source density with this function.

To verify that our dipole estimators are unbiased for the MALS sky coverage and noise structure, we create a set of simulated catalogues the same sky coverage as MALS. One simulated has the same noise everywhere, the second simulation has the σ_{20} noise of each pointing throughout the pointing, and the most realistic simulation has a varying noise level for each pointing, scaled to the σ_{20} value of that pointing. Running our dipole estimators on the simulated catalogues, we see that the dipole amplitude is underestimated for all estimates that retrieve the correct dipole direction. For estimates which are affected by incompleteness of the catalogues the dipole direction is biased toward the south pole. The completeness corrections we defined either undercompensate or overcompensate for the completeness in the simulated catalogue, however the Poisson-rms estimator correctly accounts for the noise structure in these catalogues.

While the declination estimator is able to account for the declination effect seen in the data, we are unable to constrain the declination of the dipole with this method, as it is degenerate with the cosine describing the declination systematic. As such to truly mitigate the effect it must be addressed in the processing of the data. The effects stemming from changes in (u, v) -coverage that are the main candidates of the origin of the declination effect be made mitigated by smoothing the images to a common beam. To what extent these remedy the effect remains to be seen, as other surveys are known to have similar systematics, that can in some cases persist even through the processing steps lined out.

As the dipole measurement is adversely affected by the systematic source variation in declination, the basic Poisson and Poisson-rms estimators yield only biased results. The declination estimator is able to account for this effect, however as both the declination effect and the dipole follow the same function, their parameters are degenerate and poorly constrained. Changing the prior for the declination of the dipole direction to exclude higher latitudes, we get results that are more consistent with radio dipole measured with other surveys. The result with the highest significance agrees well with the CMB dipole in terms of direction, but rejecting the CMB dipole in terms of amplitude with a significance of 5.6σ .

Outlook

Beside the measurement of the cosmic radio dipole with MeerKAT in Chapter 5, we confidently detected the dipole with 4.8σ significance with a combined measurement using RACS and NVSS (Chapter 4 and Wagenveld et al., 2023b). In addition to this, the number count dipole gained an equally significant measurement performed with AGN at infrared wavelengths (Secrest et al., 2021; Secrest et al., 2022; Dam et al., 2023). This infrared detection, which has a significance of 4.9σ , was done with data from WISE, a space-based telescope. These measurements should therefore have completely different systematics, making the fact that the dipole is detected in both cases of special significance. With the tension firmly established, the logical next step is to focus on revealing the origin of the excess dipole signal. Though we laid out some possible explanations in Chapter 1, the evidence for these is rather tenuous, and connecting these yields more of a conspiracy theory than anything else. More data will no doubt help in answering these questions, and there are interesting avenues to explore which can maximise the effectiveness of these measurements. In this chapter we will go deeper on possible avenues of future research, and what we might expect from them.

6.1 Multi-wavelength synergies

On their own, radio sources do not carry much spectral information. For a measurement of the number count dipole, we have assumed that most observed radio sources are far enough away to trace the background. From spectroscopic measurements and simulations, we know that this is indeed the case for the bulk of bright radio sources, especially AGN (e.g. Condon, 1984; Wilman et al., 2008), however for any given radio source there is some uncertainty involved. This is even more significant for MALS and other more sensitive radio surveys, which goes deep enough in flux density to probe the population of SFGs. Figure 6.1 shows the redshift distribution corresponding to a flux density distribution equal to that of the full MALS catalogue presented in Chapter 5, as taken from the SKADS catalogue (Wilman et al., 2008). This shows that many SFGs are already present in the MALS data, with many of them being appearing at low redshifts. Being able to distinguish between SFGs and AGN in the data may therefore be valuable for a dipole measurement. Though we might expect local structure to have a higher impact on a dipole measurement with SFGs, at this point it is not certain what the impact might be. In order to separate the different populations of radio sources, multi-wavelength data must be used (e.g. Algera et al., 2020; Whittam et al., 2022).




Figure 6.1: Redshift distribution of sources in the SKADS catalogue (Wilman et al., 2008) with the same flux density distribution of the full MALS catalogue.

Furthermore, recent works have made some claims that if there is evolution of spectral index, magnification bias, or the luminosity function, higher dipole amplitudes than expected can be measured, as the Ellis and Baldwin (1984) method does not take such effects into account (Dalang and Bonvin, 2022; Guandalin et al., 2022). Given these claims, a natural next step is to get redshift measurements for these sources. While spectroscopic redshifts remain time consuming to gather, the precision of photometric redshifts is ever increasing, even at higher redshifts (e.g. Duncan, 2022). These can be obtained from existing optical survey such as the Dark Energy Spectroscopic Instrument (DESI) legacy imaging surveys (Dey et al., 2019), or upcoming surveys, such as the Rubin Observatory Legacy Survey of Space and Time (LSST, LSST Science Collaboration et al., 2009). Given enough data with source classifications and photometric redshifts, the dipole could even be measured for different source populations and redshift bins, which should shed more light on the origin of the excess dipole.

6.2 Polarisation

Throughout this thesis, and indeed in most of the literature pertaining to the cosmic radio dipole, the focus is on the continuum Stokes I emission of the sources, and as a result not much is known about the dipole in polarisation. This is mostly due to the fact that less sources are available in polarisation to perform a dipole measurement with (Taylor et al., 2009). The average radio source has a polarisation fraction of 10%, which means that a measurement of a dipole requires many more sources. Nevertheless, we can make some predictions on what dipole signal we would expect to see for a population of polarised sources. Again, we can only make predictions about the kinematic part of the dipole. Tiwari and Jain (2013) specify that the kinematic dipole for the number counts of polarised




Figure 6.2: Top: Smoothed map of the fraction of sources with a polarised flux density above 0.1 mJy in NVSS. Several prominent features appear which have more polarised sources on average. These features are most likely connected to the Milky Way, as they match the features seen on the bottom plot. Bottom: De-stripped and source-subtracted Haslam et al. (1982) 408 MHz all-sky map from Remazeilles et al. (2015), showing the extent of Galactic synchrotron.

sources would have an amplitude of

$$\mathcal{D}_P = [2 + x(1 + \alpha) + x_P(1 + \alpha_P)]\beta, \quad (6.1)$$

where in addition to the variables from Equation 1.1, the power law index of the polarised flux density distribution, x_P , and the spectral index of polarised emission, α_P , are included. This is promising, as the kinematic dipole might thus be different in Stokes I and polarisation, even in the same data. This becomes more complicated however when we realise that x , α and x_P , α_P are probably not independent, so the exact dipole amplitude will be dependent on how these quantities are correlated.

So far, Tiwari and Jain (2013) is the only work that has measured a dipole in polarisation using the NVSS. They measure the dipole using polarised source counts, polarised flux weighted source counts, and degree of polarisation weighted source counts. They find a dipole with an amplitude much higher than expected from the CMB, and in fact even higher than what has been found with unpolarised emission, while a direction that agrees with the CMB dipole. Even now, NVSS is the only survey with enough polarised sources to perform a dipole estimate, though the number of usable sources in Tiwari and Jain (2013) of around 10^5 means that only high dipole amplitudes can be detected like this.

For polarisation there are even more systematic effects that can affect a measurement of the dipole, which can be not only be instrumental but physical as well. The polarised flux of a source need not be intrinsic to the source, as polarised emission can be removed or added by foreground effects such as those connected to the Milky Way. This can already be seen in NVSS, as Figure 6.2 indicates that the number counts of polarised sources are correlated to Galactic foreground features. Additionally, calibration and imaging is much more involved when dealing with polarisation. In Equation 2.13, there are terms dealing with polarisation that have to be solved in addition to the effects on Stokes I, such as polarisation leakage. These terms require observation of polarisation calibrator, which has not been done in all observations. As MeerKAT has linear polarised feeds as opposed to the more prevalent circular polarised feeds, polarisation calibration techniques are different and still under active development. Furthermore, primary beams in polarisation have a much more complex structure than in Stokes I, making it difficult to get accurate measurements off-axis. Finally, source extraction in polarisation is not straightforward, as the background noise is not generally Gaussian. All in all, a lot more work is required before polarised products are available for MALS. Even then a dipole measurement must be subject to the same or even more scrutiny than we have given the dipole in Stokes I in this thesis.

6.3 Clustering

When introducing the Poisson assumption for a distribution of radio sources in Chapter 4, an explicit choice was made to, for the moment, ignore the fact that for realistic data the true distribution would deviate from a Poisson distribution. The requirement for a Poisson point process is that all points are statistically independent, which is not the case for radio data. One effect which is highly dependent on observing frequency, angular resolution, and survey depth, is the occurrence of multi-component sources. These sources, generally AGN, have blobs of extended emission that are not connected with each other, and therefore several sources can be found by source extraction even though only a single source is present. We quantified the extent of this effect for the first ten MALS pointings in Chapter 3, showing that this effect is most prevalent at high flux densities ($S > 100$ mJy), which generally represents a very small percentage of all sources. Another effect which is physical rather than an observational effect is that of clustering of sources. Due to the way in which structure evolves in the Universe, sources tend to cluster together in overdensities. While most AGN are isotropically distributed, we see galaxy clusters containing overdensities of AGN. Going deeper, many SFGs inhabit galaxy groups, and are therefore expected to be more clustered across the board. As such, the clustering is also dependent on the depth of the survey. In MALS in particular, the effects of clustering may be exacerbated, as MALS pointings are not randomly chosen, but rather are centred on bright AGN. Such bright sources are more likely to inhabit overdensities where sources are more clustered.

These departures from a Poisson distribution can be characterised by the Poisson dispersion in

Equation 4.15. The resulting $\chi^2/d.o.f.$ value essentially gives the ratio of the variance to the mean of the distribution, which can also be interpreted (in absence of systematic effects) as the mean number of objects per cluster n_c (Peebles, 1980). We see that for a perfectly isotropic distribution, such as those generated in simulations, this value is equal to unity and the distribution is therefore perfectly Poisson. For real data however, this value is somewhat higher, and assuming all of this is due to either clustering or multi-component sources, we see that in Table 4.3 the $\chi^2/d.o.f.$ takes values of around 1-2, and values of 1.1-1.3 if we only consider the more reliable measurements.

Although the effect is not expected to mimic a dipole, taking the clustering into account may reveal additional information about the data. The way in which this may be done is presented in Siewert et al. (2020), where data from the Low Frequency Array (LOFAR) Two-metre Sky Survey (LoTSS, Shimwell et al., 2019) is tested. It is shown that rather than a Poisson distribution, a better fit is the compound Poisson distribution. The compound Poisson distribution takes the following form

$$p(n) = \frac{(N\gamma)^n e^{-N\gamma}}{n!} \frac{\lambda^N e^{-\lambda}}{N!}, \quad (6.2)$$

in which N is the amount of sources, drawn from a Poisson distribution with mean λ . The amount of observed components is n , which is drawn from a Poisson distribution with as mean the amount of sources N times the average number of components per source γ . The parameter γ does not distinguish between additional components from clustering or multi-component sources, so it describes both observed effects. In this description the Poisson dispersion is equal to $n_c = 1 + \gamma$. In terms of parameter estimation, γ can be included in the existing Poisson estimators we defined in Chapter 4 as an additional parameter to fit, which can account for the clustering of the data.

6.4 Isolating the kinematic dipole

So far, all dipole measurements in the literature have had no way of isolating the kinematic component of the dipole. Doing this might begin to answer the question of where the excess dipole originates. Some ways to separate the kinematic dipole have been proposed (Nadolny et al., 2021; Tiwari et al., 2015), but there are more possible avenues that haven't been explored. Depending on the amplitude of the respective components, these methods might require much more data than what is currently available. If for example the kinematic component would be consistent with the CMB dipole, in excess of 10^6 sources would be required to make a 3σ measurement (Crawford, 2009).

Multi-catalogue estimators

One way to distinguish between these contributions from an observational perspective requires data sets that have significantly different expected kinematic dipole amplitudes. At the current moment, the most significant test can be done between NVSS and WISE, which have both been measured by Secrest et al. (2022). Table 6.1 shows the measured dipole amplitudes along with the expected dipole from the CMB. Table 6.1 we calculate the velocity of the observer v , assuming that the dipole is entirely kinematic for both catalogues. The resulting velocities however do not agree with each other within the uncertainties. Alternatively, we can assume that the velocity is equal to the velocity measured from the CMB dipole (370 km/s), and calculate a residual dipole \mathcal{D}_{resid} . The resulting values for \mathcal{D}_{resid} for both catalogues agree within the uncertainties, favouring the interpretation that the excess

Table 6.1: NVSS and WISE dipole measurements from Secrest et al. (2022). Velocity v is calculated assuming the dipole is fully kinematic, and intrinsic dipole amplitude \mathcal{D}_{int} is calculated assuming the kinematic dipole is equal to the CMB dipole.

Catalogue	x	α	$\mathcal{D}_{kin,CMB}$ ($\times 10^{-2}$)	$\mathcal{D}_{measured}$ ($\times 10^{-2}$)	v km/s	\mathcal{D}_{resid} ($\times 10^{-2}$)
NVSS	0.78	0.75	0.42	1.23 ± 0.25	1097 ± 223	0.81 ± 0.25
WISE	1.89	1.06	0.73	1.47 ± 0.15	1310 ± 134	0.74 ± 0.15

dipole amplitude is not a kinematic effect. This is a basic measurement that can be expanded upon, for example by fitting for the dipole contributions using these catalogues in the multi-catalogue estimator introduced in Chapter 4.

Advanced estimators

The dipole estimators we defined in Chapter 4 work directly with the counts-in-cell distribution of sources, rather than the effects on the flux densities and positions of individual sources. As such, it measures the full number count dipole and does not differentiate between kinematic or other components. In Appendix A.2 we construct a likelihood based on the effects of Doppler boost, Doppler shift, and relativistic aberration occurring on individual sources. This is a basic recipe for the implementation of estimator of the cosmic dipole which could directly probe the velocity, given that we know α (for each source) and x . If successfully implemented, this can isolate the kinematic contribution to the number count dipole using a single catalogue.

Summary and conclusion

In 2019, the cosmic radio dipole was a quirk of radio data, visible in the NVSS, the most sensitive radio survey at the time. The most significant measurement of the radio dipole with the NVSS reached nearly 3σ , which in the eyes of most astronomers translates to ‘kind of weird, but not much to worry about, it’s probably a systematic we did not account for’. While other surveys were available, such as WENSS and SUMSS, none were as sensitive as the NVSS, and could at best be used in conjunction with the NVSS. TGSS yielded even higher dipole amplitudes, which has for the most part been chalked up to issues with the flux scale, strengthening the case for systematic effects being the cause of the excess dipole. Indeed, even some of those working on the cosmic radio dipole thought they would characterise these systematic effects eventually, putting the issue to rest.

Thus, starting out, the dipole was still very much unproven, so it made a great deal of sense to approach the problem from the perspective of mitigating systematic effects to make the measurement as unbiased as possible. This was the purpose of Chapter 3, where we took a deep dive into MeerKAT data from ten MALS observations. Here we checked the consistency of the calibration and imaging, automatically going through the processing logs to retrieving errors and other useful quantities such as the flux density scale. We described the creation of spectral index images, primary beam correction of the images, and the process of source finding and cataloguing. We used the rms maps created by the source finding software to characterise the noise structures in the images, finding that all pointings had similar noise structures that merely differed from each other by a constant offset, which we characterised with σ_{20} . Using mock observations with simulated data, we determined the completeness and flux recovery of the catalogues, finding that these measures could be generalised to any pointing using σ_{20} . We checked the flux density scale and astrometric accuracy by cross-matching the MALS catalogues with NVSS and RACS. For the combined catalogue of sources, we checked how extended and multi-component sources could influence a dipole measurement, and looked at the distribution of spectral indices. To see if we could correct for the incompleteness of the catalogues for a dipole measurement, we corrected differential number counts using the determined completeness, as well as with the rms coverage of each pointing. With this, we showed that we could generate a complete and homogeneous catalogue down to $300 \mu\text{Jy}$.

The next question was how to actually measure the catalogue. Many techniques used in the existing literature, for example linear estimators, were biased for incomplete sky coverage and thus would not work well on MALS data. Therefore in Chapter 4 we constructed a set of Bayesian estimators based on the assumption that counts-in-cell distribution of the data should follow a Poisson distribution.

Besides a simple estimator which just included the effect of a dipole on the number counts, we defined additional estimators which would allow for more sensitive dipole measurements. Our Poisson-rms estimator, in addition to a dipole, modelled the relation between the rms and source density, based on the assertion that incompleteness of the catalogue is directly related to the local noise. Furthermore, we defined a multi-catalogue estimator, which could estimate the dipole using multiple catalogues. We tested out these estimators on a set of simulated catalogues, as well as the NVSS and RACS catalogues, finding in all cases a dipole amplitude much higher than that expected from the CMB dipole. The most significant measurement was obtained with a joint dipole estimate on the NVSS and RACS catalogues, yielding a dipole amplitude three times higher than the CMB expectation, with a significance of 4.8σ . This is as of yet the most significant measurement of the cosmic radio dipole. We finally showed that the multi-catalogue estimator could be made to account for systematic effects in a single catalogue, by splitting the different parts of NVSS where the sky was observed with different VLA configurations, and performing a joint estimate on the separated catalogues. With this, we could alleviate the declination bias introduced by this systematic for lower flux density cuts in the NVSS catalogue.

With the complete set of 391 fully calibrated, imaged, and catalogued MALS L band observations, we obtained a catalogue covering 1632 deg^2 of sky and containing 796,916 sources. Though the noise varies throughout, the catalogue goes down to a sensitivity of $10 \text{ } \mu\text{Jy}/\text{beam}$. Combining the knowledge of the MALS and the new dipole estimator, we use this catalogue for a dipole estimate of MALS in Chapter 5. What we however found was a systematic variation of source density as a function of declination that persisted even above the completeness limit of the survey. The most plausible causes for this effect were related to the (u, v) -coverage of the observations, and characterising the most likely function the effect would follow, we defined an addition estimator which could take this into account. Before estimating the dipole on the real data however, we first ran our estimators on a set of simulated catalogues with the same sky coverage as MALS. From this we discovered that unexpectedly, the sky coverage of MALS seems to bias the estimated dipole amplitude. These simulations furthermore showed that the completeness corrections we devised are not adequately precise to recover the full set of dipole parameters. Nonetheless, quantifying these biases would allow us to correct for them on the actual data. Estimating the dipole on the full MALS dataset showed that the systematic in declination completely dominated the dipole estimates of the regular estimators. The previously defined declination estimator was however able to model the effect, but as the function it fits has the same shape as the dipole function, it was degenerate with the dipole declination, which adversely affected other parameters as well. Setting the prior for the dipole declination to exclude higher declinations improved the fit and lead to well constrained parameters. The most significant result was achieved with a $300 \text{ } \mu\text{Jy}$ flux density cut, which yielded a dipole amplitude of $\mathcal{D} = 1.57_{-0.21}^{+0.25} \times 10^{-2}$, which is four times higher than the CMB expectation. The formal significance of this measurement is 5.6σ . Although the declination could not be well constrained with the chosen estimator, its right ascension is $\text{R.A.} = 183 \pm 9$ matches the direction of the CMB dipole within 2σ . However, for a truly reliable fit, reprocessing of the data would be required to alleviate these systematics.

Even though the MALS measurement carries some caveats with it, the dipole has gained some increasingly significant measurements. In addition to our measurement of the radio dipole with a 4.8σ significance, the excess dipole has now been measured in the number counts of AGN in the infrared with a significance of 4.9σ (Secrest et al., 2021; Secrest et al., 2022). Rather than the expected discovery of the systematic that would explain it all, the excess dipole has now been measured at such different frequencies and with such different observatories that an observational systematic common

between these observations is basically impossible. This has put the discrepancy between the number count dipole and the CMB dipole nearly on par with the Hubble tension, in terms of significance and perhaps also soon in terms of independent tracers. With the dipole tension firmly established, the next steps should involve constraining the origin of the dipole. This can be done with measurements at other wavelengths, improvements of existing measurements with additional information such as polarisation, redshift, and source population, or with improved estimators.

The cosmic radio dipole is alive and well, and its implications to cosmology might very well be as important as the existing tensions, if not moreso. Once again, radio astronomy may have an important role to play in setting a new cosmological paradigm.

The radio number count dipole

A.1 Derivation

The majority of bright ($S > 1$ mJy) sources at radio wavelengths outside of the galactic plane are AGN that have a redshift distribution that peaks at $z \sim 0.8$. As such, radio sources are expected to trace the background, making it a viable probe for the cosmic dipole. Additionally, the spectral features and number count relations of AGN and other radio sources make them uniquely suitable for a dipole measurement. For most sources the dominant emission mechanism at radio wavelengths is synchrotron radiation, the spectrum of which is a power law, as described by Equation 2.1. The number density of radio sources is shown in Figure 2.4 to vary as a function of flux density, however the dipole is generally dominated by sources around the lower observed flux limit. This means that we can generally assume that the number of sources as a function of flux density $dN/d\Omega$ follows a power law relation, described by power law index x ,

$$\frac{dN}{d\Omega}(> S) \propto S^{-x}. \quad (\text{A.1})$$

The value of x can differ depending on survey depth and frequency, but usually takes values around unity. Now, we posit that we are moving with respect to the background radio sources in the sky, with some velocity $\beta = v/c$. We define a Doppler term

$$\delta = \frac{1 + \beta \cos \theta}{\sqrt{1 + \beta^2}}, \quad (\text{A.2})$$

with θ the angle between the direction of motion and some given positions on the sky, in our case radio sources. A systemic Doppler effect shifts the spectra of observed radio sources as

$$v_{obs} = v_{rest} \delta, \quad (\text{A.3})$$

which affects the flux density of these sources in the frame of the moving observer according to the spectral index α of the emission. Assuming we are observing the sky at a fixed frequency, for example

the rest frequency of the received photons, the observed flux density is shifted as

$$S_{obs}(\nu_{rest}) \propto \nu_{rest}^{-\alpha} \quad (\text{A.4})$$

$$\propto \delta^\alpha \nu_{obs}^{-\alpha} \quad (\text{A.5})$$

$$\propto \delta^\alpha S_{rest}(\nu_{obs}) \quad (\text{A.6})$$

Since the spectra of most radio sources have a negative spectral index (which is positive α), the flux density compounds with the additional effect of Doppler boosting, multiplying the observed flux density by δ . The net effect of these two terms on the observed flux density is

$$S_{obs} = \delta^{1+\alpha} S_{rest}. \quad (\text{A.7})$$

Thus given a flux limited survey of radio sources, sources are shifted above the flux density limit in the direction of motion, increasing source density, and less sources will appear in the opposite direction. In the frame of the background, this can also be interpreted as an asymmetry in the lower limit of flux density of the observer, which gets shifted as

$$S_{0,obs} = \frac{S_{0,rest}}{\delta^{1+\alpha}}. \quad (\text{A.8})$$

We translate this shift to source density following Equation A.1,

$$\left(\frac{dN}{d\Omega}\right)_{obs} = \left(\frac{dN}{d\Omega}\right)_{rest} \frac{S_{0,obs}^{-x}}{S_{0,rest}^{-x}}, \quad (\text{A.9})$$

$$= \left(\frac{dN}{d\Omega}\right)_{rest} \delta^{x(1+a)}. \quad (\text{A.10})$$

Finally, relativistic aberration caused by motion of the observer changes the positions of sources towards the direction of motion, causing a further increase in number counts in the direction of motion,

$$\tan \theta_{obs} = \frac{\sin \theta_{rest}}{\beta - \cos \theta_{rest}}. \quad (\text{A.11})$$

This, rather than having a direct effect on flux densities, changes the observed solid angle on the sky, in turn influencing source density. Going only to first order in β , the observed solid angle is shifted by

$$d\Omega_{obs} = d\Omega_{rest} \delta^{-2}. \quad (\text{A.12})$$

Though it is not unfeasible to think to measure the effects of the dipole on individual sources, the terms laid out here translate to source density in a straightforward manner with only minor (and well-justified) assumptions with regards to the ensemble of sources. To combine all these effects into a single dipole term, we make a last assumption that β is small, which is well supported by the CMB

dipole.

$$\left(\frac{dN}{d\Omega}\right)_{obs} = \left(\frac{dN}{d\Omega}\right)_{rest} \delta^{2+x(1+\alpha)}, \quad (\text{A.13})$$

$$\approx \left(\frac{dN}{d\Omega}\right)_{rest} \{1 + [2 + x(1 + \alpha)]\beta \cos \theta\}, \quad (\text{A.14})$$

$$\approx \left(\frac{dN}{d\Omega}\right)_{rest} [1 + \mathcal{D} \cos \theta]. \quad (\text{A.15})$$

Where we have now defined the dipole amplitude \mathcal{D} as

$$\mathcal{D} = [2 + x(1 + \alpha)]\beta. \quad (\text{A.16})$$

With this treatment it is straightforward to relate the observed dipole in number counts to a velocity if we know x and α . However, this is of course only valid if the observed dipole is caused by the observers' motion. As measurements of the radio dipole with large radio surveys have been found to disagree with the CMB dipole, a purely kinetic dipole should not be assumed. For this reason, the dipole is most often reported in terms of its amplitude, despite the ease with which a velocity can be derived.

A.2 Advanced dipole likelihood

Equation A.1 shows that for radio sources that number counts scale as a power law with respect to flux density. In terms of probability, this equation represents the cumulative distribution function (CDF) of the flux density of sources,

$$P(< S) = 1 - P(> S) = 1 - CS^{-x}.$$

Here C is a normalising constant to ensure that the integrated probability is equal to unity. To obtain the probability of detecting a source at a certain flux density S , we take the derivative this function,

$$\begin{aligned} P(S) &= \frac{d}{dS} (1 - CS^{-x}), \\ &= xCS^{-x-1} \end{aligned}$$

In order to normalise the probability, the resulting function must be integrated over the entire range of sampled flux densities to obtain a normalisation factor. As the power law diverges at lower flux densities, we define an absolute minimum flux density S_{min} for the distribution. This minimum flux density is defined in the rest frame, so the dipole impacts this value in the frame of the observer. The normalisation factor is computed as

$$\begin{aligned} 1 &= \int_{S_{min}}^{\infty} xCS^{-x-1} dS, \\ 1 &= CS_{min}^{-x}, \\ C &= S_{min}^x. \end{aligned}$$

The probability distribution of flux densities is then defined as

$$P(S) = \frac{xS^{-x-1}}{S_{min}^{-x}}. \quad (\text{A.17})$$

This equation defines the flux density probability distribution of sources, and in the rest frame it will be assumed that sources are uniformly distributed across the sky. From Equations A.7 and A.11, we see the dipole affects flux density and position of sources as

$$S_{obs} = \left[\frac{1 + \beta \cos \theta}{\sqrt{1 - \beta^2}} \right]^{1+\alpha} S_{rest}, \quad (\text{A.18})$$

$$\theta_{obs} = \arctan \left[\frac{\sin \theta_{rest} \sqrt{1 - \beta^2}}{\beta + \cos \theta_{rest}} \right]. \quad (\text{A.19})$$

Here θ indicates the angle between the source and the direction of motion and β is the velocity of the observer as a fraction of the speed of light.

We now take inventory of all probabilities and the related parameters. We identify our input parameters as intrinsic angular separation from dipole θ_{rest} , lower flux density limit S_{lim} , power law index of the flux density distribution x , velocity factor β , and spectral index α . The values that we measure are the noisy observed flux density \hat{S}_{obs} and noisy observed angular separation from the dipole $\hat{\theta}_{obs}$. The posterior probability from Bayes' theorem is therefore

$$P(\beta, x, \alpha, S_{lim}, \theta_{rest} | \hat{S}_{obs}, \hat{\theta}_{obs}) = \frac{\pi(\beta, x, \alpha, S_{lim}, \theta_{rest}) \mathcal{L}(\hat{S}_{obs}, \hat{\theta}_{obs} | \beta, x, \alpha, S_{lim}, \theta_{rest})}{\pi(\hat{S}_{obs}, \hat{\theta}_{obs})}. \quad (\text{A.20})$$

We are interested in working out the likelihood, which we have to convert to something that can be actually calculated.

$$\begin{aligned} \mathcal{L}(\hat{S}_{obs}, \hat{\theta}_{obs} | \beta, x, \alpha, S_{lim}, \theta_{rest}) &= P(\hat{S}_{obs} | \beta, x, \alpha, S_{lim}, \theta_{rest}) P(\hat{\theta}_{obs} | \beta, \theta_{rest}), \\ &= P(\hat{S}_{obs} | S_{obs}) P(S_{obs} | \beta, x, \alpha, S_{lim}, \theta_{rest}) \\ &\quad \times P(\hat{\theta}_{obs} | \theta_{obs}) P(\theta_{obs} | \beta, \theta_{rest}), \\ &= P(\hat{S}_{obs} | S_{obs}) P(S_{obs} | S_{rest}, \theta_{rest}, \beta, \alpha) P(S_{rest} | x, S_{lim}) \\ &\quad \times P(\hat{\theta}_{obs} | \theta_{obs}) P(\theta_{obs} | \beta, \theta_{rest}). \end{aligned}$$

These terms all represent probability distributions that we know. The lowest level terms relate the noisy observed distribution to the actual observed distribution, expected to be Gaussian

$$P(\hat{\theta}_{obs} | \theta_{obs}) \sim \mathcal{G}(\hat{\theta}_{obs} | \theta_{obs}), \quad (\text{A.21})$$

$$P(\hat{S}_{obs} | S_{obs}) \sim \mathcal{G}(\hat{S}_{obs} | S_{obs}). \quad (\text{A.22})$$

One level beyond are the functions relating the observed values to the intrinsic values through the dipole parameters already specified earlier, which can be seen as delta functions from a probability

standpoint

$$P(S_{obs}|\beta, \alpha, S_{rest}, \theta_{rest}) = \delta \left[S_{obs} - (1 + \beta \cos \theta_{rest})^{1+\alpha} S_{rest} \right], \quad (\text{A.23})$$

$$P(\theta_{obs}|\beta, \theta_{rest}) = \delta \left[\theta_{obs} - \arctan \left(\frac{\sin \theta_{rest}}{\beta - \cos \theta_{rest}} \right) \right]. \quad (\text{A.24})$$

From here only the intrinsic distributions of flux densities and positions need to be specified. The distribution of positions belongs in the prior for θ_{rest} , and is not part of the likelihood, as it is simply a uniform distribution. The distribution of flux densities follows directly from Equation A.17,

$$P(S_{rest}|x, S_{lim}) = \frac{x S_{rest}^{-x-1}}{S_{lim}^{-x}}. \quad (\text{A.25})$$

We can then implement these into the likelihood. Since S_{obs} and θ_{obs} are not observed, they are considered nuisance parameters and must be marginalised over

$$\begin{aligned} \mathcal{L}(\hat{S}_{obs}, \hat{\theta}_{obs}|\beta, x, \alpha, S_{lim}, \theta_{rest}) &= \int \int d\theta_{obs} dS_{obs} \mathcal{G}(\hat{\theta}_{obs}|\theta_{obs}) \mathcal{G}(\hat{S}_{obs}|S_{obs}) \\ &\quad \times \delta \left[S_{obs} - (1 + \beta \cos \theta_{rest})^{1+\alpha} S_{rest} \right] \\ &\quad \times \delta \left[\theta_{obs} - \arctan \left(\frac{\sin \theta_{rest}}{\beta - \cos \theta_{rest}} \right) \right] \\ &\quad \times P(S_{rest}|x, S_{lim}), \\ &= \mathcal{G} \left[\hat{\theta}_{obs} \mid \arctan \left(\frac{\sin \theta_{rest}}{\beta - \cos \theta_{rest}} \right) \right] \\ &\quad \times \mathcal{G}[\hat{S}_{obs} | (1 + \beta \cos \theta_{rest})^{1+\alpha} S_{rest}] P(S_{rest}|x, S_{lim}). \end{aligned}$$

Again, this has two more nuisance parameters, S_{rest} and θ_{rest} , that need to be marginalised over. This can however not be solved analytically, which leaves the likelihood as

$$\begin{aligned} \mathcal{L}(\hat{S}_{obs}, \hat{\theta}_{obs}|\beta, x, \alpha, S_{lim}, \theta_{rest}) &= \int \int dS_{rest} d\theta_{rest} P(S_{rest}|x, S_{lim}) \\ &\quad \times \mathcal{G}(\hat{\theta}_{obs}|\theta_{rest}, \beta) \mathcal{G}(\hat{S}_{obs}|\theta_{rest}, S_{rest}, \alpha, \beta). \end{aligned}$$

For an ensemble of sources, the likelihood is the product over all sources,

$$\begin{aligned} \mathcal{L}(\hat{S}_{obs}, \hat{\theta}_{obs}|\beta, x, \alpha, S_{lim}, \theta_{rest}) &= \prod_i^N \int \int dS_{rest} d\theta_{rest} P(S_{rest}|x, S_{lim}) \\ &\quad \times \mathcal{G}(\hat{\theta}_{obs}^i|\theta_{rest}, \beta) \mathcal{G}(\hat{S}_{obs}^i|\theta_{rest}, S_{rest}, \alpha, \beta). \end{aligned}$$

Now this likelihood can be maximised if we had all sources in the distribution starting at S_{min} . However, S_{min} is simply a convenience parameter that we defined to keep our distribution from diverging, and does not occur in nature. As long as we don't have any sources in our data which are below S_{min} in the rest frame, the likelihood will not be affected. We define S_{min} to be below

our observed lower flux density threshold, which is defined in the frame of the observer. This value, S_{thresh} , follows a dipole in the rest frame of the sources, and we only observe sources which have a value higher than S_{thresh} . The sources that are not observed must be treated as missing data, creating a truncated data set. The missing sources are accounted for by marginalising over them,

$$\begin{aligned}
 \mathcal{L}(\hat{S}_{obs}, \hat{\theta}_{obs} | \beta, x, \alpha, S_{lim}, \theta_{rest}) &= \prod_i^N \int \int \mathcal{G}(\hat{\theta}_{obs}^i | \theta_{rest}, \beta) \mathcal{G}(\hat{S}_{obs}^i | \theta_{rest}, S_{rest}, \alpha, \beta) \\
 &\quad \times P(S_{rest} | x, S_{lim}) dS_{rest} d\theta_{rest}, \\
 &= \prod_i^{N-m} \int \int dS_{rest} d\theta_{rest} P(S_{rest} | x, S_{lim}) \\
 &\quad \times \mathcal{G}(\hat{\theta}_{obs}^i | \theta_{rest}, \beta) \mathcal{G}(\hat{S}_{obs}^i | \theta_{rest}, S_{rest}, \alpha, \beta) \\
 &\quad \times \prod_i^m \int_0^{2\pi} \int_{S_{lim}}^{S_{thresh}} d\theta_{mis} dS_{mis} P(S_{rest} | x, S_{lim}) \\
 &\quad \times \mathcal{G}(\hat{\theta}_{mis}^i | \theta_{rest}, \beta) \mathcal{G}(\hat{S}_{mis}^i | \theta_{rest}, S_{rest}, \alpha, \beta).
 \end{aligned}$$

Implementing this likelihood computationally turned out to be quite a challenge and had to be abandoned due to time constraints. Though the integrals can be done numerically, integrating over the missing data requires one to know how many sources are exactly missing, or having the number of sources as an additional nuisance parameter. Whether this implementation is feasible or even efficient remains to be seen.

APPENDIX B

Paper I

The paper *Wagenveld et al. (2023a)*, *A&A*, 673, A113, 30 was published in *Astronomy & Astrophysics* under the Creative Commons CC-BY 4.0 license, and is reproduced in its original form below.

The MeerKAT Absorption Line Survey: Homogeneous continuum catalogues towards a measurement of the cosmic radio dipole[★]

J. D. Wagenveld¹, H.-R. Klöckner¹, N. Gupta², P. P. Deka², P. Jagannathan³, S. Sekhar^{4,3},
S. A. Balashev^{5,6}, E. Boettcher^{7,8,9}, F. Combes¹⁰, K. L. Emig¹¹, M. Hilton^{12,13}, G. I. G. Józsa^{1,14},
P. Kamphuis¹⁵, D. Y. Klutse¹², K. Knowles¹⁴, J.-K. Krogager¹⁵, A. Mohapatra², E. Momjian³,
K. Moodley¹², S. Muller¹⁶, P. Petitjean¹⁷, P. Salas¹⁸, S. Sikhosana¹², and R. Srianand²

¹ Max-Planck Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany
e-mail: wagenveld@mpi-fr-bonn.mpg.de

² Inter-University Centre for Astronomy and Astrophysics, Post Bag 4, Ganeshkhind, Pune 411 007, India

³ National Radio Astronomy Observatory, Socorro, NM 87801, USA

⁴ The Inter-University Institute for Data Intensive Astronomy, Department of Astronomy, and University of Cape Town, Private Bag X3, Rondebosch 7701, South Africa

⁵ Ioffe Institute, 26 Politekhnicheskaya st., St. Petersburg 194021, Russia

⁶ HSE University, Saint Petersburg, Russia

⁷ Department of Astronomy, University of Maryland, College park, MD 20742, USA

⁸ X-ray Astrophysics Laboratory, NASA/GSFC, Greenbelt, MD 20771, USA

⁹ Center for Research and Exploration in Space Science and Technology, NASA/GSFC, Greenbelt, MD 20771, USA

¹⁰ Observatoire de Paris, Collège de France, PSL University, Sorbonne University, CNRS, LERMA, Paris, France

¹¹ Jansky Fellow of the National Radio Astronomy Observatory, 520 Edgemont Road, Charlottesville, VA 22903, USA

¹² Astrophysics Research Centre and School of Mathematics, Statistics and Computer Science, University of KwaZulu-Natal, Durban 4041, South Africa

¹³ Wits Centre for Astrophysics, School of Physics, University of the Witwatersrand, 2050 Johannesburg, South Africa

¹⁴ Department of Physics and Electronics, Rhodes University, PO Box 94 Makhanda 6140, South Africa

¹⁵ Ruhr University Bochum, Faculty of Physics and Astronomy, Astronomical Institute, 44780 Bochum, Germany

¹⁶ Department of Space, Earth and Environment, Chalmers University of Technology, 43992 Onsala, Sweden

¹⁷ Institut d'astrophysique de Paris, UMR 7095, CNRS-SU, 98bis bd Arago, 75014 Paris, France

¹⁸ Green Bank Observatory, Green Bank, WV 24944, USA

Received 16 November 2022 / Accepted 16 February 2023

ABSTRACT

The number counts of homogeneous samples of radio sources are a tried and true method of probing the large-scale structure of the Universe, as most radio sources outside the Galactic plane are at cosmological distances. As such, they are expected to trace the cosmic radio dipole, an anisotropy analogous to the dipole seen in the cosmic microwave background (CMB). Results have shown that although the cosmic radio dipole matches the direction of the CMB dipole, it has a significantly larger amplitude. This unexplained result challenges our assumption of the Universe being isotropic, which can have large repercussions for the current cosmological paradigm. Though significant measurements have been made, sensitivity to the radio dipole is generally hampered by systematic effects that can cause large biases in the measurement. Here we assess these systematics with data from the MeerKAT Absorption Line Survey (MALS), a blind search for absorption lines with pointings centred on bright radio sources. With the sensitivity and field of view of MeerKAT, thousands of sources are observed in each pointing, allowing for the possibility of measuring the cosmic radio dipole given enough pointings. We present the analysis of ten MALS pointings, focusing on systematic effects that could lead to an inhomogeneous catalogue. We describe the calibration and creation of full band continuum images and catalogues, producing a combined catalogue containing 16 307 sources and covering 37.5 square degrees of sky down to a sensitivity of 10 $\mu\text{Jy beam}^{-1}$. We measure the completeness, purity, and flux recovery statistics for these catalogues using simulated data. We investigate different source populations in the catalogues by looking at flux densities and spectral indices and how they might influence source counts. Using the noise characteristics of the pointings, we find global measures that can be used to correct for the incompleteness of the catalogue, producing corrected number counts down to 100–200 μJy . We show that we can homogenise the catalogues and properly account for systematic effects. We determine that we can measure the dipole to 3σ significance with 100 MALS pointings.

Key words. surveys – galaxies: statistics – radio continuum: galaxies

1. Introduction

The vast majority of sources seen at radio wavelengths outside of the Galactic plane are known to be at cosmologically

significant distances ($\langle z \rangle \sim 0.8$; e.g. Longair & Scheuer 1966; Condon & Ransom 2016). This makes homogeneous samples of radio sources ideal for studying the local luminosity function, along with the large-scale structure and evolution of the Universe (Longair & Scheuer 1966). The number counts of radio sources were used as evidence against a static Euclidean Universe (Ryle & Scheuer 1955), providing a convincing argument

[★] The catalogue is only available at the CDS via anonymous ftp to [cdsarc.cds.unistra.fr](ftp://cdsarc.cds.unistra.fr) (130.79.128.5) or via <https://cdsarc.cds.unistra.fr/viz-bin/cat/J/A+A/673/A113>

in favour of a strongly evolving universe even before the discovery of the cosmic microwave background (CMB; [Penzias & Wilson 1965](#)). As radio sources trace the large-scale structure of the Universe, they are expected to abide by the cosmological principle, which asserts that the Universe is homogeneous and isotropic. However, there is an anisotropy expected in the number counts of radio sources, caused by the velocity of the Solar System with respect to the cosmological background. This expresses itself as a dipole and is the dominant anisotropy observed in the CMB ([Lineweaver 1997](#)). The movement of the observer induces Doppler boosting and relativistic aberration that cause the apparent luminosity and position of radio sources to shift, resulting in a dipole in the number counts of radio sources. A measurement for the cosmic radio dipole was first proposed by [Ellis & Baldwin \(1984\)](#), who showed that 2×10^5 sources, adequately distributed along the axis of the dipole, are required for a 3σ measurement of the radio dipole, assuming the Solar System velocity derived from CMB measurements.

Using data from the National Radio Astronomy Observatory (NRAO) VLA Sky Survey (NVSS; [Condon et al. 1998](#)), [Blake & Wall \(2002\)](#) made the first significant measurement of the dipole, with a direction and amplitude that correspond to those of the CMB. Subsequent studies were performed with the NVSS and other radio surveys, such as the Westerbork Northern Sky Survey (WENSS; [Rengelink et al. 1997](#)), the Sydney University Molonglo Sky Survey (SUMSS; [Mauch et al. 2003](#)), and the Tata Institute for Fundamental Research (TIFR) Giant Metrewave Radio Telescope (GMRT) Sky Survey's first alternative data release (TGSS ADR1, [Intema et al. 2017](#)). It was found that the cosmic radio dipole, while being consistent in terms of direction with that of the CMB, significantly differs from the CMB dipole in terms of amplitude (e.g. [Singal 2011](#); [Rubart & Schwarz 2013](#); [Tiwari & Jain 2013](#); [Tiwari et al. 2015](#); [Colin et al. 2017](#)). These early dipole measurements found that survey-wide systematic effects, which cause varying source densities, can greatly bias dipole estimates. This is usually remedied by strict cuts in flux density, which dramatically decrease the number of usable sources. Even with these flux density cuts, some surveys, such as TGSS, yield anomalous dipole results that have been attributed to systematics due to problems with flux calibration (e.g. [Singal 2019](#); [Bengaly et al. 2018](#); [Siewert et al. 2021](#)). While results differ depending on the survey and estimator used, the amplitude of the radio dipole is consistently larger (by a factor of 2–6) than the amplitude of the CMB dipole (see [Siewert et al. 2021](#), for an overview), while the direction of the dipole remains consistent. With similar results found using the number counts of active galactic nuclei (AGN) at infrared wavelengths ([Secrest et al. 2021, 2022](#); [Singal 2021](#)), it becomes increasingly difficult to explain them with systematic effects or faulty analysis. Only in a recent analysis by [Darling \(2022\)](#) was a dipole found to be consistent with the CMB in both direction and amplitude; this was done by combining the VLA Sky Survey (VLASS; [Lacy et al. 2020](#)) and the Rapid Australian Square Kilometre Array Pathfinder (ASKAP) Continuum Survey (RACS; [McConnell et al. 2020](#)), though it presents only one counterpoint to the many works that find an increased dipole amplitude. Considering a purely kinematic origin of the dipole, the cosmic radio dipole and the CMB dipole are in obvious tension with each other. The excess dipole found in the radio therefore must be a result of a different process, which could have major implications for cosmology. As radio galaxies trace the underlying matter distribution, a dipole in their distribution would break with

isotropy, one of the fundamental assumptions of cosmology. The assumption of isotropy and homogeneity is founded on the notion that we as observers do not occupy a special place in the Universe, and these results suggest that there is some flaw in this assessment.

Working towards an independent measurement of the radio dipole, we utilise the MeerKAT Absorption Line Survey (MALS; [Gupta et al. 2016](#)), a deep radio survey with pointings centred on bright radio sources. MALS is carrying out a dust-unbiased search for neutral hydrogen (HI, 21 cm) and hydroxyl (OH, 18 cm) absorption lines at redshifts $0 < z < 2$ in order to unravel the processes driving the steep evolution of the star formation rate density. As a blind search for absorption lines, every MALS pointing is centred on a bright AGN (>200 mJy at 1.4 GHz). The targets have been chosen from the NVSS and SUMSS catalogues and are cross-checked with Wide-field Infrared Survey Explorer (WISE) data in order to build a dust-unbiased sample of AGN ([Gupta et al. 2022](#)). Early results show that MALS is able to attain unprecedented sensitivity to absorption lines in these bright AGN ([Gupta et al. 2021](#); [Combes et al. 2021](#)). In addition to the search for absorption lines, the data taken will be sensitive enough to produce deep continuum images, down to $10 \mu\text{Jy beam}^{-1}$. With a full width at half maximum (FWHM) field of view of 1 degree in the L band (1.27 GHz), each MeerKAT pointing presents a few square degrees and potentially thousands of sources. With 391 pointings currently observed in the L band, the full survey will provide thousands of square degrees of deep continuum sky and hundreds of thousands of sources.

Though the expected MALS number counts are sufficient for a dipole measurement, a dipole estimate requires a homogeneous catalogue. Systematic effects influencing the sensitivity of surveys are common and are usually dealt with by making conservative cuts in the data to avoid biasing the dipole estimate. Instead, in this work we present a thorough analysis of ten MALS pointings, aiming to fully understand the systematics present in the survey data. This will allow us to account for these systematics when measuring the radio dipole using hundreds of MALS pointings. The nature of the survey provides additional challenges for this type of measurement. Previously, measurements of the dipole have been performed with contiguous surveys such as the NVSS, whereas MALS will be sparser, sampling the sky in many different directions. However, compared to these surveys, MeerKAT has a much higher sensitivity ($10 \mu\text{Jy beam}^{-1}$), which allows us to probe deeper into the population of faint radio sources. Furthermore, past dipole measurements from contiguous sky surveys have been performed post-factum, with little knowledge of the internal processing, and therefore present systematics of these surveys beyond what is described in the literature. In this paper we study the first ten continuum images of MALS in depth in order to assess their quality. We investigate the systematics in calibration, imaging, and source finding on image quality and source counts, and extrapolate our findings to the rest of the survey.

This paper is organised as follows. In Sect. 2 we describe the MALS data. The initial creation of the source catalogues and the completeness measures are described in Sect. 3. In Sect. 4 we describe results from the full catalogue of sources. We investigate how different source populations affect the catalogues in Sect. 5. In Sect. 6 we assess the prospects for a dipole measurement with MALS using the results in this paper. Finally, in Sect. 7, we summarise the findings of this paper.

Table 1. Calibration details of the pointings presented in this paper, grouped by observation runs.

Flux cal	Flux density 870 MHz (Jy)	Target	Gain cal	Flux density 1365 MHz (Jy)	Reference flux ^(a) 1400 MHz (Jy)	Spectral index ^(a)	Distance from target (degrees)
J1331+3030 J1939-6342	18.829 14.095	J2023-3655	J2052-3640	1.45 ± 0.005	1.367 ± 0.017	-1.258 ± 0.02	5.8
J0408-6545 J1939-6342	27.027 14.095	J0126+1420	J0108+0134	3.24 ± 0.01	3.113 ± 0.070	-0.273 ± 0.02	13.5
J1331+3030 J1939-6342	18.829 14.095	J1133+0015 J1232-0224 J1312-2026	J1150-0023 J1256-0547 J1311-2216	2.86 ± 0.005 10.66 ± 0.05 5.5 ± 0.01	2.9 ^(b) 9.82 ± 0.120 4.857 ± 0.060	- -0.490 ± 0.05 -1.281 ± 0.04	4.2 6.9 1.8
J0408-6545 J1939-6342	27.027 14.095	J0001-1540 J0006+1728	J2357-1125 J2253+1608	2.12 ± 0.009 15.5 ± 0.02	1.8 ^(b) 16.199 ± 0.198	- -0.193 ± 0.03	4.4 17.5
J0408-6545 J1939-6342	27.027 14.095	J0240+0957 J0249-0759 J0249+0440	J0238+1636 J0240-2309 J0323+0534	0.61 ± 0.002 6.15 ± 0.003 2.85 ± 0.002	0.528 ± 0.014 5.938 ± 0.131 2.766 ± 0.062	-0.246 ± 0.05 -0.154 ± 0.03 -0.920 ± 0.01	6.7 31.2 8.5

Notes. The flux densities of the calibrators are reported by the CASA fluxscale task, which changes reference frequency based on whether the source is a flux calibrator or gain calibrator. ^(a)<https://skaafrika.atlassian.net/wiki/spaces/ESDKB/pages/1452146701/L-band+gain+calibrators> for properties of MeerKAT *L*-band calibrators. ^(b)Value from the old list of calibrators, no longer publicly available.




Fig. 1. Sky distribution of the first 391 observed pointings of MALS. The Galactic plane is largely avoided, and since 89% of the pointings are selected directly from NVSS, the vast majority of pointings are above a declination of -40 degrees. The pointings used in this analysis are highlighted in red.

2. MALS data

The distribution of the first 391 observed pointings of MALS is shown in Fig. 1. In order to assess the data quality of the individual MALS pointings and the impact for the dipole estimates an initial set of ten pointings, shown in Fig. 1 in red, has been selected out of five observing runs to probe different ranges of right ascension, declination, and central source flux density.

2.1. Observations and calibration

The general setup of a single MALS observation includes observations of three science targets and corresponding calibrators. The observation is scheduled with a flux calibrator observed for 10 min at the start and end of each observing run. Each target is observed for 20 min at a time, cycling through all targets three times for a total observing time of an hour per source. Before and after each target observation, a nearby gain calibrator is observed for one minute. Cycling between targets like this maximises the UV-coverage with minimal increase in overhead. Observing

multiple targets in a single run is not only convenient in terms of processing, but is also critical in taking stock of systematic effects, such as flux density scale or phase errors, potentially introduced during observation or calibration. All observations have a correlator integration time of 8 s, with observations carried out in 32 K mode, providing 32 768 channels with a channel width of 26.123 kHz. With a frequency range of 856–1712 MHz, the total bandwidth is 856 MHz, with a central frequency of 1.285 GHz.

The MeerKAT data are shipped to the Inter-University Centre for Astronomy and Astrophysics (IUCAA) in India and processed by the Automated Radio Telescope Imaging Pipeline (ARTIP). The complete deployment of ARTIP in MALS is described in Gupta et al. (2021). ARTIP presents an environment where data can be processed according to user specifications and is based on the Common Astronomy Software Applications (CASA) tasks (The CASA Team et al. 2022). Each dataset undergoes a round of basic flagging, removing known radio frequency interference (RFI) frequencies. This is followed by flux calibration, bandpass calibration, and gain calibration, each step having the possibility of additional automated flagging. The final target visibilities used for the imaging process are produced by applying the flags and calibration solutions.

As part of the overall evaluation of the individual pointings, all the available information was assessed automatically with an evaluation scheme that has been developed to trace errors of the calibration process by searching through the logging information of ARTIP. This scheme also extracts relevant information from the logs, such as the flux densities of the calibrator sources. An overview of the targets and calibrators of the ten selected pointings, organised by observation block, is shown in Table 1. For the gain calibrators, both the flux density determined during calibration and from a reference catalogue is listed.

2.2. Self-calibration and continuum imaging

For the purposes of continuum imaging, the data are averaged over 32 channels and divided into 16 spectral windows (SPWs), resulting in 64 channels per SPW. Once again, frequencies with

Table 2. Details on all ten pointings after complete processing and source finding.

Target	RA	Dec	Flux density	Flux density	Spectral index	PSF maj	PSF min	PSF PA	Counts	σ_{20}	Demerit
			NVSS (mJy)	MALS (mJy)		"	"	°			
J0001-1540	00 ^h 01 ^m 41 ^s .57	−15°40′40″.60	436	513.6 ± 1.0	−0.85 ± 0.04	7.7	6.3	−8.9	2132	26	14.5
J0006+1728	00 ^h 06 ^m 47 ^s .35	+17°28′15″.40	226	220.1 ± 0.7	−0.28 ± 0.08	11.4	6.3	−7.1	1378	29	10.1
J0126+1420	01 ^h 26 ^m 13 ^s .24	+14°20′13″.10	577	685.6 ± 0.7	−0.95 ± 0.09	10.5	6.3	−2.5	1591	33	8.6
J0240+0957	02 ^h 40 ^m 27 ^s .19	+09°57′13″.00	521	589.3 ± 0.7	−1.11 ± 0.09	10.2	6.6	9.1	986	48	18.7
J0249-0759	02 ^h 49 ^m 35 ^s .41	−07°59′21″.00	646	711.0 ± 0.6	−0.97 ± 0.09	9.2	6.6	−1	2619	19	7.7
J0249+0440	02 ^h 49 ^m 39 ^s .93	+04°40′28″.90	420	472.6 ± 0.3	−0.80 ± 0.09	8.1	6.7	−7.6	1558	29	6.3
J1133+0015	11 ^h 33 ^m 03 ^s .12	+00°15′48″.90	233	377.7 ± 0.9	−0.01 ± 0.07	8.9	6.7	−15.3	803	52	15.9
J1232-0224	12 ^h 32 ^m 00 ^s .13	−02°24′04″.10	1647	1823.4 ± 5.2	−0.31 ± 0.09	8.5	6.7	−9	611	73	20.8
J1312-2026	13 ^h 12 ^m 07 ^s .86	−20°26′52″.40	778	851.4 ± 0.4	−0.83 ± 0.09	7.7	6.3	−9.1	2431	21	12.2
J2023-3655	20 ^h 23 ^m 46 ^s .21	−36°55′21″.20	436	406.6 ± 0.5	−0.11 ± 0.12	10.9	6.8	−88	2160	22	5.8

known strong RFI are flagged (see also Fig. 2 of Gupta et al. 2021). The resulting dataset has a total of 960 channels, a bandwidth of 802.5 MHz (869.3–1671.8 MHz), and a central frequency of 1.27 GHz as a result of the edges of the band being flagged.

As each field contains a strong point source at its centre, both phase and amplitude self-calibration can be performed (Cornwell & Fomalont 1989). In total, three phase and one amplitude calibration steps are performed, with imaging each step to improve the local sky model. As is common with self-calibration in CASA, we use the clean components created in `tclean` as the local sky model for calibration. We iterate on the model by creating masks for `tclean` using the Python Blob Detection and Source Finder (PYBDSF; Mohan & Rafferty 2015). Starting with a mask containing only the central source, after a set number of iterations PYBDSF is used on the image to create the mask for the full field, initially with a high S/N threshold and lowering the threshold for subsequent runs to gradually expand the model. Creating the clean masks in such a way ensures that cleaning is mostly limited to real emission, while also speeding up the imaging by limiting the cleaning area.

Though the self-calibration can be a significant improvement on the image it can also be potentially unstable. To monitor the stability of solutions, a diagnostic tool for self-calibration produces a report on the variation of relevant statistics such as noise and central source flux density in different steps of calibration. As with calibration, the logs were evaluated for errors and warnings during the self-calibration process and information relevant to assessing the quality of the products, such as percentage of flagged data and theoretical noise limit, were extracted.

Imaging is performed using Multi-term Multi-Frequency Synthesis (MTMFS; Rau & Cornwell 2011) deconvolution with four pixel scales (0, 2, 3, and 5 pixels) to model extended emission in the images and two Taylor terms to account for the spectral shape of the sources. This produces two Taylor term images, which describe the spectral shape of the emission to zeroth and first order, respectively. As such, the zeroth order Taylor term I_0 represents the continuum flux density of the field at the reference frequency of 1.27 GHz, while the first order Taylor term I_1 describes the spectral index,

$$I_0 = I_{\nu_0}^{\text{sky}}; I_1 = \alpha I_{\nu_0}^{\text{sky}}. \quad (1)$$

To maintain a balance between sensitivity and resolution in the images, visibilities are weighted using Briggs weighting (Briggs 1995) with robust value of 0. Because we are imaging with a large field of view, we use W-projection (Cornwell et al. 2005) with 128 projection planes to correct for the fact

that our baselines are non-coplanar. The final data products consist of the restored, model, residual, sum-of-weights, and point spread function (PSF) images for both Taylor terms. Furthermore, spectral index, spectral index error, and mask images are also produced. The continuum images have a pixel size of 2'' and a size of 6000 × 6000 pixels. This results in a square image of 3.3 degrees on a side. Though individual pointings have different beams, as detailed in Table 2, they are on average aligned in the north-south direction, with a mean major axis of 9.3'' and mean minor axis of 6.5''.

2.3. Spectral index images

The L band of MeerKAT has a bandwidth of 802.5 MHz, which is large enough to be sensitive to the spectral shape of the radio emission within the band. If this is not taken into account when imaging the full band, this incurs a large uncertainty in flux density. The general solution for this is MTMFS deconvolution, which models the frequency dependence of the emission with a Taylor expansion. In our case, as mentioned in Sect. 2.2, we model the frequency dependence of the emission in the pointings to first order in ν . With this we can create maps describing the spectral index α , defined by the relation between flux density S and frequency $S \propto \nu^\alpha$, of the emission in the image.

Although MTMFS imaging also produces a spectral index image, pixels below 5 times the peak residual are masked in this image. To retain flexibility, we therefore chose to produce the spectral index images from the Taylor term images ourselves. From the definition of the Taylor term images in Eq. (1), a spectral index image can be obtained using $\alpha = I_1/I_0$, from which we will be able to measure the spectral indices of sources. To keep values in the spectral index image from diverging, pixels are masked where values in the Stokes I image are below 10 $\mu\text{Jy beam}^{-1}$. When measuring the spectral index in some region of the image, usually defined by the extent of a source, we assign a spectral index as the intensity weighted mean of the measured pixels in the spectral index image, with intensity weighted standard deviation as the error,

$$\bar{\alpha} = \frac{\sum_{i=1}^n I_{0,i} \alpha_i}{\sum_{i=1}^n I_{0,i}}, \quad (2)$$

$$\sigma_\alpha = \sqrt{\frac{\sum_{i=1}^n I_{0,i} (\alpha_i - \bar{\alpha})^2}{\frac{n-1}{n} \sum_{i=1}^n I_{0,i}}}. \quad (3)$$

If more than half of the measured pixels in a region are invalid in the spectral index image, this carries over to the measured spectral index and uncertainty by assigning a masked value.

2.4. Primary beam correction

Due to the primary beam response of the MeerKAT antennas, sources away from the pointing centre appear fainter than they are in reality. As this effect is not corrected for in the imaging stage, resulting continuum images will have accurate flux densities at the pointing centre but attenuated flux densities that become fainter the farther from the pointing centre they are located. A simplified model of the primary beam is described in [Mauch et al. \(2020\)](#), which assumes the primary beam of MeerKAT as directionally symmetric, describing it with a cosine-tapered illumination function,

$$P(\rho, \theta_{\text{pb}}) = \left[\frac{\cos(1.189\pi\rho/\theta_{\text{pb}})}{1 - 4(1.189\rho/\theta_{\text{pb}})^2} \right]^2. \quad (4)$$

Here ρ is the distance from the pointing centre and θ_{pb} represents the angular size of the FWHM of the primary beam, a quantity that is dependent on the observing frequency, ν ,

$$\theta_{\text{pb}}(\nu) = 57'.5 \left(\frac{\nu}{1.5 \text{ GHz}} \right)^{-1}. \quad (5)$$

At the central frequency of our continuum images of 1.27 GHz, the FWHM of the primary beam is $\theta_{\text{pb}} = 67'$. This simplified model is implemented in the `katbeam`¹ PYTHON package. As the primary beam is frequency dependent, it affects the spectral index images, increasing the measured spectral index away from the pointing centre. The spectral index change induced by the primary beam can be approximated by

$$P_\alpha(\rho, \nu) = -8 \log(2) \left(\frac{\rho}{\theta_{\text{pb}}} \right)^2 \left(\frac{\nu}{\nu_0} \right)^2, \quad (6)$$

Again, we assume the frequency ν to be equal to the central frequency $\nu_0 = 1.27$ GHz.

In reality, the MeerKAT primary beam in the L band is more complicated and cannot be completely described by a directionally symmetric model. [de Villiers & Cotton \(2022\)](#) present and analyse holographic measurements of the MeerKAT primary beam, showing the directional asymmetries present due to variations between individual antennas. For an accurate model of the primary beam, we use these holographic measurements to correct our images. As we utilise the full 802.5 MHz bandwidth of the L band for these images, a primary beam correction must take this into account. Though a wideband primary beam correction is implemented in the CASA task `widebandpbcor`, there are no models of the MeerKAT beams available. As such, we implement the wideband primary beam correction ourselves using the same basic recipe, which consists of creating a primary beam with a frequency structure matching that of the image, in this case creating a primary beam model for each of the 16 SPWs of the continuum data. As in the imaging step, we model the multi-frequency primary beam with two Taylor terms. The primary beam corrected Taylor term images are then defined as follows:

$$I'_0 = P_0^{-1} I_0, \quad (7a)$$

$$I'_1 = P_0^{-1} \left(I_1 - \frac{P_1}{P_0} I_0 \right). \quad (7b)$$

Here, P_0 and P_1 represent the zeroth and first order Taylor term primary beams, respectively, where P_0/P_1 should be equal to α_{pb} as specified in Eq. (6).




Fig. 2. Flux offsets of the central sources and gain calibrators per field. Gain calibrators are compared to their reference flux density (circles) as specified in Table 1 and with their NVSS counterparts (diamonds). Central sources are compared to their NVSS counterparts (squares), and NVSS flux densities are all converted to rest frequency assuming $\alpha = -0.75$. Sources are ordered by observing run, separated by vertical black lines.

While we use the holographic wideband primary beam corrections described in Eq. (7) for the main results of this work, we also briefly explore the simplified corrections of Eqs. (4) and (6) and see how they compare to the wideband corrections. At applying the primary beam corrections, the image is cut off at the 5% level of the primary beam (at the central frequency of 1.27 GHz), which leaves us with a circular image with a diameter of approximately 4000 pixels, or 2.2 degrees. As a result of reduced sensitivity towards the edges of the image, the noise is increased there.

2.5. Assessment of calibration

Processing the raw data to the final scientific data products can introduce errors, affecting the flux density scale. A first order estimation of the flux density scale can be obtained by comparing the flux densities of the gain calibrators with their literature values. As discussed in Sect. 2.1, we evaluate the automated process of the data calibration by generating diagnostic reports and automatically evaluating logged information in order to determine problems in the data processing. This evaluation singles out errors and warnings present in the logs, allowing direct insight into any problems that might have occurred during the calibration process. Furthermore, it extracts information we can use to assess the quality of calibration from the logs, such as the flux densities of calibrator sources.

Table 1 summarises the observation and calibration details, showing the targets and their associated calibrator sources. The flux densities of the flux- and gain calibrators are extracted from the logs and the flux densities of gain calibrators are compared to the MeerKAT reference catalogue ([Taylor & Legodi 2021](#)). We extend this to a broader assessment of the flux density scale in Fig. 2, where we show the flux density offsets of the gain calibrators and central sources of the individual pointings. Along with the comparison in Table 1, both gain calibrators and central sources (see Table 2) are compared to their NVSS counterparts. Flux densities are corrected for frequency using the spectral index from the reference catalogue if

¹ <https://github.com/ska-sa/katbeam>

available, assuming $\alpha = -0.75$ otherwise. Combining the measurements from the ten pointings, the mean flux density ratios are 1.03 ± 0.26 between the gain calibrators and their NVSS counterparts, 1.07 ± 0.07 between gain calibrators and their reference values, and 1.08 ± 0.19 between central sources and their NVSS counterparts. We note that the absolute amplitude calibration of NVSS is based on Baars et al. (1977) and has an uncertainty of up to 12% with respect to the here used Perley & Butler (2017) scale, depending on the calibrator used.

We note that the SUMSS and NVSS measurements were taken with different instruments at different times, so some variation is to be expected. The current assessment does not include astrometric precision, as calibrators are not imaged. We assess this aspect along with the another flux density scale assessment by cross-matching the full catalogue of sources with other surveys in Sect. 4.2.

2.6. Assessment of image quality

With any radio image, there is a great number of variables that can influence the quality of the image, both related to intrinsic properties of the pointing and to the process of calibration and self-calibration. As discussed in Sect. 2.2, a report is generated that monitors image statistics such as noise and central source flux during different self-calibration steps. Furthermore, the logs are automatically evaluated for possible errors and warnings and information relevant to the quality of the self-calibration and imaging is extracted. To evaluate the final image product, the image quality of the individual ten pointings is assessed by using the root mean square (RMS) noise maps that are automatically produced during the source finding procedure by PyBDSF (see Sect. 3). In particular, we investigate the overall noise characteristics by evaluating the sky coverage with respect to the RMS noise. A direct measurement of the noise allows us then to easily correlate image quality with other characteristics of the pointings.

We create a smoothed representation of the ten pointings by median stacking their normalised RMS noise images, which is shown in Fig. 3. As all pointings have a strong source at their centre, the noise is increased at the pointing centre and increases towards the edges of the pointing as a consequence of the primary beam response. Figure 3 shows that some directional effects are left in the image. Notably, there is an elongated noise structure in the centre, associated with the bright central source, aligned in the north-south direction. The stacked beam included in the figure aligns well with the elongated structure, indicating that the most prominent structures are a result of the shape of the stacked PSF of the images. The imprint of the stacked PSF is also the most likely cause of the cross-like structure seen in the stacked image. Though we have the wide-band primary beam correction based on holographic images that take into account the asymmetries present in the primary beam, pointings are observed for three separate blocks of 20 min in an observing night, which smears out the asymmetries in the primary beam². This effect cannot be easily corrected for in the image plane, but could be taken into account during imaging using A-projection (Bhatnagar et al. 2008). Though present, the asymmetries here are small and dominated by the other noise structures in the image.

The usual method of determining RMS noise in an image relies on measuring RMS noise in an area close enough to the




Fig. 3. Median stacked pixel values of the RMS noise images of all ten pointings. As primary beam correction is applied, the noise goes up towards the edges of the image. Since a strong central source is always present, the noise is always higher in the centre as well. The given contours from white to black are 20, 40, and 60% RMS noise coverage. The stacked beam (50 time increase in size) of all the pointings is shown in the lower-left corner, matching the elongated structure in the centre.




Fig. 4. RMS noise coverage for all ten pointings. The dotted line indicates the 20% coverage level, which is used to define σ_{20} . Noise varies appreciably between the pointings; however, the overall structure of the RMS noise coverage curves remains consistent, indicating that σ_{20} is a good zeroth order measure of the noise scale.

pointing centre to not be affected by the primary beam and far enough from strong sources to be unaffected by artefacts. Due to the number and structure of MALS pointings, this cannot be reliably done in an automated fashion. Instead, we investigate the differences in noise level between individual pointings by using the RMS noise images to assess the RMS noise coverage, measuring the cumulative distribution of noise levels across an image. Figure 4 shows the RMS noise coverage curves for the individual fields, and it can be seen that RMS noise coverage curves are similar in structure but offset from each other. To quantify this offset and thus characterise the noise in the individual pointings, we define σ_{20} at 20% RMS noise coverage, representing the noise in the central portion of the image (see Fig. 3). We will see that σ_{20} excellently serves as a normalisation factor to account for the differences in noise levels between the

² MeerKAT antennas have Alt-Az mounts, such that the sky rotates with respect to the dish while observing.

pointings, and can be used to unify the assessment of individual pointings and extend them to the full survey.

There are several factors which can contribute to the overall noise level in an image, not all of which are easily quantifiable. However, an important aspect to consider is the shape of the synthesised beam or PSF, determined by the UV-coverage of the observation, which in turn is determined by the array configuration, observing time, and elevation of the target at the time of observation. There are two aspects to the PSF that influence image noise. A measurement of RMS noise in Jy beam^{-1} will be influenced by the shape of the beam³, and very bright sources can have persistent and bright sidelobes from the shape of the PSF that are difficult to clean completely and as a result push up the noise in an image. To quantify this last effect we calculate the demerit score detailed in [Mauch et al. \(2020\)](#) to estimate the contributions of bright sources to RMS noise in the image. We calculated the independent source contributions to the errors in the image using all sources that have an unattenuated flux density of more than 100 mJy. The demerit score, d , is then defined as

$$d = \left[\sum_{S > 100 \text{ mJy}}^i \left\{ \left(\frac{8 \ln(2) \rho \sigma_p}{\theta_{pb}^2} + \sigma_g \right) S_{a,i} \right\}^2 \right]^{1/2}, \quad (8)$$

where the first term represents the contribution of pointing error σ_p scaling with distance from the pointing centre ρ , and the second term is the receiver gain error σ_g . The contribution of each source comes in the form of their attenuated flux density $S_{a,i}$. Appropriate values for the MeerKAT L band are detailed in [Mauch et al. \(2020\)](#), which we also use ($\theta_{pb} = 67'$, $\sigma_p = 30''$, $\sigma_g = 0.01$). The demerit scores of all pointings are included in [Table 2](#). A correlation is present between demerit score and σ_{20} , and especially pointings with high σ_{20} show increased demerit scores. Though pointings with lower σ_{20} show more scatter in their demerit scores, this nonetheless shows demerit score as a first order estimate of pointing quality, which we can utilise as a predictive measure.

3. Source finding

With thousands of sources expected to be detected in every MALS pointing, we require an automated source finding algorithm to find and characterise these sources. A small number of these are suitable for radio images, and perform comparatively similar ([Hale et al. 2019](#)). Of these, PYBDSF has been used in several recent data releases of large-scale surveys, such as the LOFAR Two-Metre Sky Survey (LoTSS; [Shimwell et al. 2019](#)) and the RACS ([Hale et al. 2021](#)). PYBDSF stands out in its ability to model extended emission with its wavelet decomposition module, and provides easy ways to compile source catalogues and assess the quality of the fields. Besides generating catalogues, PYBDSF provides output maps related to the input image, such as the RMS noise images we used in [Sect. 2.6](#), and mean and residual images. Once RMS noise and mean maps are obtained PYBDSF allows these maps to be used as input to ensure source finding is performed with the exact same parameters. For MALS, we thus make use of PYBDSF, both for creating clean masks during self-calibration as detailed in [Sect. 2.2](#), and integrating PYBDSF into the workflow to automatically carry out source finding, cataloguing,

cross-matching and combining catalogues, using PYTHON-based scripts developed by the authors⁴.

In order to understand the impact of the individual pointings to a general catalogue, we evaluate the source finding procedure for each pointing. We investigate completeness (what fraction of sources do we detect) and purity (what fraction of sources is real) in source counts with respect to signal-to-noise ratio, flux density, and source size, as well as PYBDSF's capability to accurately recover flux densities.

3.1. Stokes I catalogues

In order to compile a source catalogue from PYBDSF, various steps that depend on the initial setup of PYBDSF are needed. PYBDSF identifies islands of emission that are brighter than the island threshold. Within these islands PYBDSF finds emission peaks above a corresponding pixel threshold, and for each peak found fits a 2D Gaussian to the peak and surrounding emission. Performing source finding on our MALS images, we impose an island threshold of 3σ , and a pixel threshold of 5σ . Individual Gaussian components are combined into sources in a way that can be specified by the user, and we elect to combine Gaussian components that occupy the same island into a single source. The RMS noise in the images is determined by a sliding box, and we decrease the size of the sliding RMS box near bright sources to avoid spurious detections of artefacts around these sources as much as possible. Furthermore, to improve fitting of extended sources in the field, we enable à trous wavelet decomposition ([Holschneider et al. 1989](#)). The PYBDSF settings can be summarised as follows:

```
thresh_isl = 3.0
thresh_pix = 5.0
rms_box = (150, 30)
adaptive_rms_box = True
adaptive_tresh = 100
rms_box_bright = (40, 15)
group_by_isl = True
atrous_do = True
atrous_orig_isl = True
atrous_jmax = 3.
```

For the purposes of analysing and building the final catalogue, we required a number of output products from PYBDSF. The output from source finding includes both a catalogue of sources and of individual Gaussian components. Furthermore, a background RMS noise and background mean image are produced, as well as a residual image.

For a single pointing, we ran source finding using PYBDSF and modified the output source catalogues by adding to the existing columns the ID of the MALS pointing, a source name following IAU convention, and the distance of the source to the pointing centre. As PYBDSF does not calculate spectral indices unless an image has multiple channels, we measured the spectral index of the sources in our fields from the spectral index images as described in [Sect. 2.3](#), using the extent of the Gaussian (major axis, minor axis, position angle) of the source to define the region in the image.

Though PYBDSF is configured to avoid spurious detections as much as possible, it is unavoidable that some artefacts are falsely identified as sources. We identified artefacts around the ten brightest sources in each image by flagging sources within five times the major axis of the beam that have less than 10% of

³ The clean beam of an image is determined during imaging by fitting a 2D Gaussian to the central lobe of the PSF.

⁴ <https://github.com/JonahDW/Image-processing>

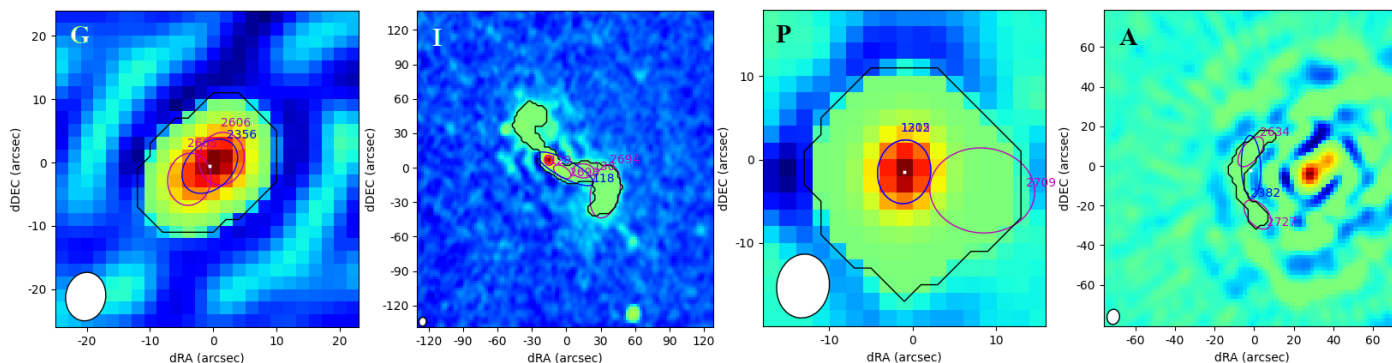


Fig. 5. Examples of possible source classes. The black outline shows the island threshold, the magenta ellipses show the individual Gaussian components fit to the source, and the blue ellipses show the combined Gaussian describing the source. From left to right: (i) an elongated source fit by two Gaussian components. The combined Gaussian describes the source adequately, and it has been assigned the ‘G’ class. (ii) A likely FRI source with complex structure, better described by the island than the Gaussian components. It has been assigned the ‘I’ class. (iii) A point source with an additional noise peak that has been fitted with a Gaussian component. It has been assigned the ‘P’ class. (iv) An artefact caused by a nearby bright source. It has been assigned the ‘A’ class.

the peak flux density of the bright source. This is largely motivated by the shape of the PSF, which can have sidelobes with a strength of up to 10% of the maximum. Though this does not get rid of all false detections in the image (see Sect. 3.2.4), it flags the most prominent imaging artefacts.

To assess the quality of the Gaussian fitting by PYBDSF, we performed a visual inspection on select sources. We created cutouts from the images and performed visual inspection, which was implemented in a separate module based on PYTHON and CASA⁵. PYBDSF assigns each source a flag indicating whether the source is fit by a single Gaussian (‘S’), multiple Gaussian components (‘M’), or Gaussian component(s) on an island with other sources (‘C’). Since all Gaussian components that occupy the same island are always combined into one source, the ‘C’ flag is not present in our catalogues. For the visual inspection, we considered all sources made up of multiple Gaussian components. As such, all sources that carry the ‘M’ flag – which make up 8% of the all sources found in the fields – are flagged for visual inspection. Through the visual inspection, we then assigned an additional flag indicating the nature of the source and how well it is described by the PYBDSF model:

G: sources that are well described by the Gaussian model.

I: complex sources that are not adequately described by the Gaussian components fit to them. The flux density of these sources is better described by the integrated flux of the island, and their position by the flux weighted mean position of the island.

P: sources fit with multiple Gaussian components where only one is required to adequately describe the source. Other Gaussian components are likely fit to noise fluctuations coinciding with the source.

A: artefacts that will be flagged as such in the catalogue.

Figure 5 shows an example for each of the cutout classes, and how we identify the different possible cases. To aid in visual inspection, in the source cutout we plot the individual Gaussian components, the combined source Gaussian, and the island threshold. Therefore, in this step we use both the source catalogues and the Gaussian component catalogues. Along with the cutout classes, additional columns are added to the table that describe the sources. In the cutouts we measure the integrated flux density of the island, the spectral index of the island, the




Fig. 6. Assigned flags to sources with multiple Gaussians, separated by the number of Gaussian components fit to them by PYBDSF. Keeping in mind that most of these sources are fit with two Gaussian components, the ‘I’ class is preferred for sources with three or more components, while the ‘P’ class consists almost exclusively of sources with two components.

intensity weighted mean position of the source, and a flag indicating if these measures are valid. Additionally, the number of Gaussian components is recorded for each source, as the initial PYBDSF catalogue only indicates whether a source has been fit with multiple Gaussian components or not.

Figure 6 shows the classification of all sources in the ten pointings that have been fit with multiple Gaussian components, 1259 in total. We see that almost all 120 sources assigned with the ‘P’ class have two Gaussian components, and a relatively large percentage of sources with the ‘I’ class have more than two Gaussian components assigned. Around 185 (15%) of these sources were considered to be adequately described by their Gaussian components, while 946 (75%) are more complex and better described by their island attributes. Only 8 sources are flagged as obvious artefacts.

3.2. Evaluation of individual pointings

In order to determine the reliability of the source finding routine and to assess how detection of sources is affected by their properties, we measured the completeness, purity, and flux recovery

⁵ <https://github.com/JonahDW/CASA-Poststamp>




Fig. 7. Completeness for unresolved sources (left) and resolved sources (right) for the different fields and their associated uncertainties as a function of flux density. There are large differences between the pointings: pointings with higher noise levels have lower overall completeness. The completeness is lower for resolved sources as well.

statistics of the catalogues. Here we assess these qualities for individual pointings to see how characteristics of the pointings such as central source flux density and noise level affect these quantities.

To measure completeness and flux recovery, we required complete knowledge of the intrinsic flux densities and shapes of the sources that are present in the image. To that end, we used realistic samples of simulated extragalactic radio sources from the Wilman et al. (2008) simulation of the SKA Design Study (SKADS). Though more recent simulations such as the Tiered Radio Extragalactic Continuum Simulation (T-RECS; Bonaldi et al. 2019) are available, the SKADS catalogues include morphology details of all sources and source components, which is necessary information when injecting sources into the data. From the SKADS simulations we created mock catalogues with 5000 sources uniformly distributed in flux density that have a flux density above $10 \mu\text{Jy}$, which is equal to the limit of thermal noise ($10 \mu\text{Jy beam}^{-1}$) for an unresolved source. With this we allowed for the possibility of noise fluctuations to push sources above the detection threshold. We injected sources from the mock catalogue uniformly distributed into the residual images produced by PYBDSF, which are devoid of sources but share the noise characteristics of the original images. We then performed the source finding routine again on these images, using the same mean and RMS noise maps determined by PYBDSF from the original image. This ensures that source finding is performed in the exact same way as the original image. We considered a source recovered if it is detected within the FWHM of the major axis of the clean beam from the original position. In order to reach a more robust measure, this process was repeated 50 times for each pointing, separately for point sources and resolved sources.

3.2.1. Completeness

With the procedure described above, we can make a statistically robust assessment of the completeness in the pointings. The (source) completeness in this case simply gives the fraction of sources that is detected, most commonly measured as a function of flux density of the source. The completeness curves for the individual pointings, for both resolved and unresolved

sources, can be appreciated in Fig. 7. Not only is there a large difference between resolved sources and unresolved sources, pointings individually have large differences between them as well. To investigate other aspects of the completeness, we look at the fields J0249-0759 and J1232-0224, which have the lowest and highest noise levels among the pointings, respectively (see Table 2), which should yield the most extreme cases and allow us to probe variation between the fields.

Unresolved sources. As the SKADS catalogues describe the intrinsic shapes of sources, we can assess completeness for point sources by only injecting sources with a major axis of zero. The sources are defined in the image as delta functions, and convolved with the clean beam of the individual image. As the total flux density of point sources is concentrated in one peak, they are much easier to detect relative to resolved sources. Point sources allow us to assess completeness without being affected by source morphology, and so we use them to determine the completeness with respect to distance from the pointing centre. As sensitivity decreases outwards from the centre, we expect completeness to decrease as well. Figure 8 shows source completeness as a function of both flux density and distance from the pointing centre for the pointings J0249-0759 and J1232-0224. It is clear that indeed the completeness decreases with increased distance from the pointing centre, but is also lower near the pointing centre. This is a direct result of the strong source at the centre of each pointing pushing up the noise in its immediate vicinity. In the case of J1232-0224, which has a very strong source at the pointing centre, there is significant impact on completeness in the central portion of the image. To investigate the relation between the completeness and noise floor as a function of distance to the pointing centre we use the RMS maps created by PYBDSF. By radially averaging the RMS noise of the pointing, we obtain RMS noise as a function of distance to the pointing centre. As we have set the detection threshold at 5σ , we plot the radially averaged 5σ detection curve in Fig. 8, showing that this curve almost perfectly follows the completeness ‘transition zone’ for both pointings. In this transition zone the completeness goes up steeply from zero to one, and the flux density at which this occurs is directly related to the noise floor.




Fig. 8. Completeness for unresolved sources as a function of flux density and distance from the pointing centre, ρ , for the fields J0249-0759 (left) and J1232-0224 (right). The radially averaged 5σ curves (black lines) for the corresponding pointings are seen to follow the zone where completeness transitions from zero to one. Due to the presence of a strong central source in J1232-0224, completeness is lower in the central region of this pointing. Pixels with no sources in them have been coloured grey.

Resolved sources. We performed the same experiment for resolved sources, where we define a resolved source as a source that has major axis and minor axis larger than 0 in the SKADS catalogues⁶. Sources are randomly selected out of the catalogue, so the distribution of source shapes injected in the image represents the distribution of the SKADS sample. These sources are injected as Gaussians into the image, and as with point sources, convolved with the clean beam. Owing to their lower surface brightness, resolved sources are often less easily detectable compared to point sources with the same flux density. To check how the size of sources affects completeness, we define the area ratio Q_A of a source as the ratio between the area of the source and the beam as defined by their Gaussian characteristics. These are the major and minor axes θ_{maj} and θ_{min} for the source and B_{maj} and B_{min} for the clean beam of the pointing,

$$Q_A = \frac{\theta_{\text{maj}}\theta_{\text{min}}}{B_{\text{maj}}B_{\text{min}}}. \quad (9)$$

We show the completeness as a function of area ratio and flux density in Fig. 9 for the pointings J0249-0759 and J1232-0224. As shown in Fig. 8, the completeness of unresolved sources is related to the noise floor. This same relation should be present for resolved sources, in addition to the relation between completeness and source size. In order to disentangle the two different contributions to completeness for resolved sources, we divided flux densities by the ratio of local noise to the lowest noise in the image. The result in Fig. 9 shows completeness for uniform noise, so that only the source size and flux density affect completeness. We see a power law decrease (linear in the log-log scale) in completeness as a function of area ratio in both pointings, indicating that this is a universal feature for our source detection. This can be easily understood by considering that for larger sources the total flux density is divided over a larger area, which decreases the peak flux density that is used to detect these sources.

⁶ A small subset of sources with minor axis of 0 and major axis larger than zero are not included in the simulations.

3.2.2. Flux recovery

Using the mock catalogues, we investigate the ability of PYBDSF to accurately recover flux densities. This can be checked by looking at the flux densities measured by PYBDSF relative to the input flux densities from the mock catalogues. This is an important quality to verify as deviations from the expected 1:1 relationship are obviously undesirable. In Fig. 10 we show the measured flux densities against input flux densities of the pointings J0249-0759 and J1232-0224. We see that on average sources have a flux density that matches with their input value. There is however a portion of sources with lower input flux density that have a significantly higher measured flux density than their input. These sources have their flux densities boosted by noise fluctuations, which are present in various orders of strength in the images, from thermal noise to calibration artefacts. We expect these sources to land on positive as well as negative noise peaks, but only sources on positive peaks will be detected. This results in an Eddington bias (Eddington 1913) pushing up the distribution of flux densities. To make quantitative statements about this bias, we need to combine data from all the 10 pointings, which we do in Sect. 4.3.

3.2.3. Limitations of simulations

The method we have used here for measuring completeness and flux recovery relies on injecting sources directly into the residual images and measuring their properties with PYBDSF. The advantage of this is a direct probe into the machinery of PYBDSF, as this is the only ‘black box’ between the input sources and the measurement. However, these simulations ignore some effects that affect the flux densities and shapes of sources in radio data, such as calibration effects, clean bias, and averaging effects like time and bandwidth smearing (Bridle & Schwab 1999). Probing these requires injecting sources directly into the visibilities and reprocessing the image, something that is not efficient for a large survey such as MALS. Finally, sources are injected into the image convolved with the clean beam as opposed to the PSF. This is well motivated for brighter sources,




Fig. 9. Completeness for resolved sources as a function of flux density and ratio between the area of the source and the beam Q_A , for the fields J0249-0759 (left) and J1232-0224 (right). Completeness can be seen to linearly decrease in the log-log scale as a function of Q_A , showing that larger sources are harder to detect. Flux densities are compensated for the local noise in order to equalise completeness for different positions in the image. Pixels with no sources in them have been coloured grey.




Fig. 10. Input flux density plotted against the measured flux density, for the fields J0249-0759 (left) and J1232-0224 (right), based on 100 simulations. As the noise floor in J1232-0224 is relatively high ($\sigma_{20} = 80 \mu\text{Jy beam}^{-1}$), only sources above $200 \mu\text{Jy}$ are detected, while in J0249-0759 sources are detected down to $50 \mu\text{Jy}$.

as these have been mostly cleaned during the imaging process. For faint sources this is not the case, especially since the masks for cleaning are generated by PYBDSF and are thus subject to the same selection that we used for the final images. To make the simulations more realistic, all undetected sources should therefore be convolved with the PSF. It is not clear how this should affect source finding, but the general consequence of this is that below the detection threshold sources immediately become fainter as a consequence of being convolved with the PSF rather than the clean beam. The PSF also spreads the emission of these sources over a large area, which could affect RMS noise if source crowding is high enough. This would however only be the case if

images would be close to or at the confusion limit, which is not the case for MeerKAT in the L band down to at least $0.25 \mu\text{Jy}$ (Mauch et al. 2020).

3.2.4. Purity

The purity, or inversely the false detection rate, measures what fraction of the sources detected in the image are true detections. For a well chosen detection threshold, the amount of false detections in an image is expected to be small. It is important to have a handle on the amount of false detections, as it should be taken into account when calculating number counts. To determine




Fig. 11. Purity of catalogues for different pointings. The fraction of false detections is indicated by the red line. The dashed red line indicates the fraction of false detections without flagging artefacts. The open histograms show the number of sources detected in the pointings, with the filled histograms indicating the number of false detections.

the purity, we invert the pixel values of the images and run PYBDSF on the inverted images, using the RMS and inverted mean maps determined by PYBDSF from the original image. This again ensures that source finding is performed in the same way as on the original image. Since all real sources have positive flux density the only sources detected in the inverted images will be false detections. These false detections broadly fall into three categories, which we differentiate as noise peaks, (calibration) artefacts, and ghost sources. Noise peaks are statistical outliers of noise and can therefore appear at any point in the image, and are symmetric around the mean, such that these sources detected in the inverted image correspond roughly to the false detections in the normal image. Artefacts are sidelobes found around strong sources in the image, making them more easily traceable. As described in Sect. 3.1, we consider a source to be an artefact if they are found within 5 times the major axis of the clean beam of the ten brightest sources in the image, and have less than 10% of the peak flux density of the bright source. As the brightest negative sidelobe of the PSF is in general twice as bright as the brightest positive sidelobe, we would expect more artefacts to be found in the negative image. This seems consistent with the data, as using this criterion for artefacts flags 44 of the 241 sources found in the inverted images, while flagging 22 sources in the pointing catalogues. Finally, there are ghost sources, which appear as negative sources too bright to be noise fluctuations, in some cases even strong enough that they have sidelobes that are detected as sources. These sources can be caused by calibration with an incomplete sky model, and only have faint positive counterparts (Grobler et al. 2014). Strong ghosts can add to the number of false detections with their sidelobes, but only a handful of such cases are seen in the images.

We plot the amount of false detections per pointing in Fig. 11, both in terms of absolute counts (coloured bars) and fraction (red line). The amount of false detections strongly depends on pointing, and we find two pointings that are most strongly affected: J0001-1540 and J1232-0224. The latter is affected by a strong central source, which leads to reduced number counts and an increased fraction of false detections, while the former would be considered a good pointing, both in terms of number counts and noise properties. There does appear

however a cluster of relatively bright (10–100 mJy) sources present far from the pointing centre, which can contribute to noise. The presence of a number of strong sources far out in the field has also potentially affected self-calibration, as a high number of ghost sources are seen in the image. This result suggests that purity of any individual pointing is not always easily predictable, and should each be assessed separately.

4. Combined catalogue

The combined source catalogue of ten pointings contains 16 307 sources, and covers 35.7 square degrees of sky. In the previous section we mostly assess the quality of individual pointings. Here we combine the catalogues of the individual pointings to increase statistical power, which allows us to investigate subtler systematic effects that affect all pointings.

4.1. Correcting residual primary beam effects

In Sect. 2.4 we described primary beam corrections to both the flux densities and spectral indices in the images. Besides the main wideband primary beam corrections using holographic images, we also described corrections with simplified analytic forms. Before investigating the difference between these methods, we must make additional corrections to residual primary beam effects. In general, the simplified analytical corrections work well up to the FWHM of the primary beam, but farther out results begin to diverge. This is an effect that is seen in both the spectral indices as well as the flux densities of sources, mainly caused by using the primary beam correction based on the central frequency $\nu_0 = 1.27$ GHz for the entire bandwidth of 802.5 MHz. In order to take into account the contribution of the entire bandwidth, we recalculate the corrections by integrating over the bandwidth rather than assuming the frequency to be equal to ν_0 . The necessary corrections are computed for each source in the catalogue separately, depending on distance from the pointing centre and spectral shape.

In the images corrected by the simplified analytical function of Eq. (4), flux densities of sources appear higher the farther they are from the pointing centre. In order to properly correct for the effect in flux density, the spectral index of the source must be known or assumed. For reasons explored in Sect. 5.2, we cannot trust all spectral indices to be accurate and perform this correction assuming $\alpha = -0.75$ for all sources. We then calculate a correction factor for the flux densities as a function of distance from the pointing centre. The assumed primary beam model is as before (Eq. (4)), and the correction is computed by integrating the primary beam over the frequency range of the band. Considering a source with some spectral index α , the flux density of the source is described by $S(\nu, \alpha) \propto \nu^\alpha$. Due to the effect of the primary beam, the flux density of the source has some attenuation factor $a(\rho, \alpha)$ applied to it. This factor is described by the primary beam:

$$a(\rho, \alpha) = \frac{\int_{\Delta\nu} S(\nu, \alpha) P(\rho, \nu) d\nu}{\int_{\Delta\nu} S(\nu, \alpha) P(0, \nu) d\nu}. \quad (10)$$

Since the flux densities have already undergone primary beam correction, we need to correct for the ratio between this term and the correction from Eq. (10),

$$S_{\text{corr}} = \frac{P(\rho, \nu_0)}{a(\rho, \alpha)} S_{\text{measured}}. \quad (11)$$

The effect of this correction should become visible when comparing flux densities to external catalogues. If the flux densities are properly corrected, the flux density ratio between catalogues should be constant as a function of distance to the pointing centre.

In contrast to the flux densities, both the analytical function from Eq. (6) and the wideband primary beam correction from Eq. (7) leave a residual effect in the spectral indices of sources farther from the pointing centre. Taking Eq. (6) to describe the spectral index induced by the primary beam variation, we correct these values taking the full bandwidth into account, recalculating the effect of the primary beam on spectral indices by integrating over the bandwidth, $\Delta\nu$:

$$P_{\alpha,\text{int}}(\rho) = -8 \log(2) \left(\frac{\rho}{\theta_{\text{pb}}} \right)^2 \int_{\Delta\nu} \left(\frac{\nu}{\nu_0} \right)^2 d\nu. \quad (12)$$

To correct the already measured spectral indices present in the catalogues we subtract the difference between the integrated and original primary beam correction term,

$$\alpha_{\text{corr}} = \alpha_{\text{measured}} + [P_{\alpha,\text{int}}(\rho) - P_{\alpha}(\rho, \nu_0)]. \quad (13)$$

4.2. Cross-matching catalogues

To further investigate systematics that affect the pointings on a more general level, we continue the assessment from Sect. 2.5, now using the sources of the entire field. We cross check our sources with their counterparts from NVSS and RACS, as all our pointings here are within the sky coverage of these two surveys. Cross-matching was performed by checking whether source ellipses, defined by the 3σ extent of the Gaussians describing these sources, overlap between the catalogues. We required a minimum overlap in area of 80% to consider sources to be a match. Sources in one catalogue could be matched with any number of sources in the other, to account for different resolutions between the catalogues. Due to uncertainties in position and flux density near the NVSS detection threshold of 2.5 mJy, sources below 5 mJy in NVSS were not considered. We find that 997 sources are matched to NVSS, of which 845 are matched to a single source, and 2064 sources are matched to RACS, with 1949 matched to a single source.

There are a number of factors that can influence astrometric precision of an observation, such as errors in the reference frequency or timestamps. Some of these errors were present in earlier MeerKAT observations (e.g. Mauch et al. 2020; de Villiers & Cotton 2022). While the errors should no longer be an issue, it is important to cross check positions in the field with an external catalogue for potential astrometric errors. The astrometric offsets of sources to their NVSS counterparts can be seen in Fig. 12, where the offsets are shown only for single matched sources. Overall, the offsets are very small with a median offset of $\sim 0.3''$, which is less than one-sixth of the image pixel size ($2''$) and well within the uncertainty. The scatter in both directions is smaller than $3''$, which is less than the semi-minor axis of the average clean beam of $3.25''$, also shown on the figure.

In Sect. 4.1 we corrected spectral indices and flux densities accounting for residual effects introduced by the frequency range covered by the band. Cross-matching sources with external catalogues is an important check of the correctness of their measured flux densities. Figure 13 shows the flux density ratio of 845 MALS and NVSS sources, and Fig. 14 shows the flux density ratio of 1949 MALS and RACS sources. Only sources that are matched to a single source are used, and flux densities




Fig. 12. Astrometric offsets to NVSS for all ten pointings combined. The median offsets are given by the grey dashed lines, with the grey area indicating the uncertainty. The majority of sources lie within a FWHM of the average minor axis of the clean beam. The data are binned where five or more sources occupy the defined bin area; otherwise, individual sources are shown.




Fig. 13. Ratio of flux densities of the sources in MALS compared to their NVSS counterparts as a function of distance from the pointing centre (ρ). The running median flux density ratio of the analytical primary beam correction both with (blue line) and without (red line) the corrections made in Sect. 4.1 are shown, as well as the running median flux density ratio of the holographic wideband primary beam correction (green line).

have been converted to the MALS rest frequency (1.27 GHz) assuming $\alpha = -0.75$. In both figures the corrections for the residual primary beam effect have properly re-scaled the flux densities, as the median flux density ratio (blue line) stays largely consistent with distance from the pointing centre, but a residual effect is left towards the edges of the image. In Fig. 13 it stands out immediately that there is a systematic offset between the MALS and NVSS flux densities, as MALS flux densities are 18% higher on average. The overall flux density scale offset is $S_{\text{MALS}}/S_{\text{NVSS}} = 1.18 \pm 0.26$. In contrast to NVSS, Fig. 14 shows that the flux densities between MALS and RACS agree extremely well up to $\rho \sim 0.5$ degrees. The overall flux density scale offset is $S_{\text{MALS}}/S_{\text{RACS}} = 1.06 \pm 0.39$, with the 6% overall offset originating mostly from the outer parts of the primary beam.




Fig. 14. Ratio of flux densities of the sources in MALS compared to their RACS counterparts as a function of distance from the pointing centre (ρ). The running median flux density ratio of the analytical primary beam correction both with (blue line) and without (red line) the corrections made in Sect. 4.1 are shown, as well as the running median flux density ratio of the holographic wideband primary beam correction (green line).

Though the result from NVSS might indicate any number of problems that could cause the offset, the additional data from RACS rules out most of these assumptions. A likely source of uncertainty would be the assumption of spectral index; however, this would impact the RACS results far more significantly, with its rest frequency of 887 MHz. From the RACS flux density offset we can assume that the 6% offset stems mostly from the residual primary beam effect, but this can only explain part of the offset seen in NVSS. If this offset is persistent, it points to a systematic effect affecting either NVSS or both MALS and RACS. Due to the relatively low sensitivity of the surveys, only about 10% of MALS sources are matched to a counterpart, which makes the error bars on the flux density offset measurement rather large. As such, the measured offset is within the uncertainty, preventing us from making any definitive statement on the flux density offset. Combined with the measurement from Sect. 2.5, the flux density scale of MALS does not currently significantly deviate from the expected value, but the offset seen here indicates that more data are needed.

4.3. General assessment of the complete catalogue

In Sect. 3 the individual pointings have been evaluated with respect to completeness, purity, and flux density recovery. Here these properties are assessed on the entire catalogue in order to understand the impact of these characteristics on the final data product.

4.3.1. Completeness

To statistically determine completeness of the data, we refactored the completeness to make it consistent between different pointings. In order to achieve this, instead of expressing completeness as a function of flux density in units of Jansky, we show flux density in units of σ_{20} as defined in Sect. 2.6. As we showed in Sect. 3.2.1, completeness for point sources scales linearly with local noise, and thus should be 0.2 at $5\sigma_{20}$ for all pointings. Though σ_{20} has units of Jy beam^{-1} , for point sources integrated flux density and peak flux density are in principle equal, which

means in this definition σ_{20} has the same value in Jy. Figure 15 shows that in terms of σ_{20} , pointings have very similar completeness curves, which allows us to combine the individual pointings and evaluate completeness for the whole survey, as indicated by the black combined completeness curve.

Combined completeness is also assessed as a function of separation from the pointing centre using only point sources, and as a function of major axis of the source using resolved sources. Both are shown in Fig. 16. Combining the completeness from all the pointings gives enough statistical power to paint a clear picture of how the completeness is dependent on these variables. A clear relation is shown between completeness and distance from the pointing centre. The major difference between individual pointings seems to be the influence of the central source on the completeness. These differences are however extremely well modelled by the radially averaged RMS noise (see Fig. 8), indicating that completeness is related to the local noise. The right plot in Fig. 16 shows that there is a power law decrease in completeness for larger sources. This was already suggested in Fig. 9, but with the combined catalogues we have enough number counts to fully cover the space.

4.3.2. Flux recovery

In evaluating the individual pointings in Sect. 3.2, faint sources on average had higher measured flux densities compared to input flux densities. To fully assess this effect, we combine the flux recovery statistics from all pointings in Fig. 17. In the combined statistics the effect is clearer, with bins farther away from the flux density ratio of unity being occupied with on average 10 sources per bin. There is no visible dependence on flux density or distance from the flux density ratio of unity. Assuming a Poisson distribution of these bins with mean and variance $\lambda = 10$, we take all bins with fewer than 25 sources (5σ) to be part of this distribution. These bins combined contain 1.7% of all sources, indicating that this effect is rather small in terms of induced bias. Up to this point we have assumed that the flux density measured from the Gaussian fitting (the `Total_flux` column in the catalogues) best represents the flux density of the sources. In Fig. 17 we compare the flux recovery between the Gaussian flux density and the integrated flux density from the island that the source occupies, where the contour indicates the threshold of 100 sources per bin. We see that across the board Gaussian flux densities are skewed towards higher values, where island flux densities remain symmetric around the input flux density. This is an effect that can significantly affect our catalogues, especially considering the increased number counts at lower flux densities. Consequently, in the analysis in the rest of the paper, we assume that the flux densities of sources are more accurately represented by the island flux density.

4.3.3. Purity

Combining false detections from all ten pointings, 241 sources are detected in the inverted images, making up 1.5% of the combined catalogue. As described in Sect. 3.2.4, our artefact identification criterion flags 44 of these, leaving 197 sources, or 1.2% of the combined catalogue. With the combined catalogue of false detections, we can investigate how purity is affected by other variables such as flux density and distance from the pointing centre. This will allow us to properly account for the purity of the catalogue, in order to not overestimate number counts. Given the variety of ‘source types’ seen in negative images and what counterparts we expect to see in the positive, the overall




Fig. 15. Completeness as a function of flux density for unresolved (left) and resolved (right) sources for the different fields, re-factored with σ_{20} and combined (black curves). Re-factoring the completeness curves to σ_{20} shows clearly that they are simply shifted with respect to each other, and we can define a unified completeness measure for the survey as a function of σ_{20} for both resolved and unresolved sources.




Fig. 16. Combined source completeness as a function of distance from the pointing centre (unresolved sources, left) and major axis of the source (resolved sources, right). The left plot reflects the overall structure of the pointings, and shows that completeness is quite straightforwardly a radially averaged version of the noise structure as shown in Fig. 3. Note that the flux density is normalised by σ_{20} . The right plot indicates a clear power law relation between the size of sources and the completeness, where larger sources are on average less complete.

amount of false detections should be lower than the amount of sources seen in the negative image. In this sense, the purity is more appropriately an upper limit rather than a direct measure of false detections.

Figure 18 shows the combined purity as a function of flux density and distance from the pointing centre. The left plot shows the purity as a function of flux density and shows that the fraction of false detections increases with higher flux density. This is largely a result of the overall number of sources decreasing at higher flux density, but does show that the flux density distribution of false detections is not the same as that of real sources. The lack of false detections at low flux densities shows that our 5σ detection threshold does not lead to a lot of spurious detections. It is noteworthy that more sources are flagged as artefacts at higher flux densities, indicating that artefacts around bright sources have higher flux densities on average. The right plot of Fig. 18 shows a strong dependence of purity on the

separation from the pointing centre, similar to the completeness. False detections increase near the central source because of strong artefacts, and there is a steady increase towards the edges of the primary beam. We see that our artefact selection criterion correctly picked out most of the artefacts originating from sidelobes of the central source, which dramatically increases the purity in the central portion of the image. Although the number of false detections restricts the statistical power of these results, the relations already show clear trends for the purity as a function of flux density and distance from the pointing centre that can be used when assessing number counts.

4.4. Resolved and unresolved sources

For the analysis of completeness in Sects. 3.2.1 and 4.3.1, we assumed that our catalogues are populated with both unresolved or resolved sources, and that these should be assessed separately.




Fig. 17. Input flux density plotted against the measured flux density for both Gaussian flux densities (left) and Island flux densities (right). The threshold of 100 sources per bin (black contour) shows quite clearly the bias present in Gaussian flux densities compared to island flux densities.




Fig. 18. Purity of catalogues as a function flux density (left) and separation from the pointing centre (right). The fraction of false detections is indicated by the red line, both the total fraction (dotted line) and with removal of sidelobes (solid line). The open histograms show the number of sources detected in the pointings, with the filled histograms indicating the number of false detections. Though there seems to be no strong relation between flux density and purity, the number of false detections is strongly dependent on distance from the pointing centre, increasing both towards the centre and towards the edges of the pointing. Our criterion for identifying artefacts flags most of the false detections around the central source.

Figure 7 shows this distinction is warranted, as these source types have very different completeness relations. If we want to apply this knowledge to real sources in the catalogue, we must have a reliable way of determining whether a source is resolved. We expect sources are resolved when their size exceeds the size of the synthesised beam of the image; however, we must take the uncertainties introduced by noise in the image and fitting errors into account.

We determine source size by measuring the ratio between integrated flux density S and peak flux density S_{peak} of the source, which should be equal to one for an unresolved source. Figure 19 shows S/S_{peak} as a function of S/N , for both resolved and unresolved sources in our combined catalogue. Here and in subsequent usage of S/N we define it as the ratio between the peak flux density of the source and the local RMS. Due to a combination of uncertainties, unresolved sources follow a log-normal distribution in S/S_{peak} (Franzen et al. 2015), and thus a normal distribution in $R = \ln(S/S_{\text{peak}})$ with mean 0 and standard




Fig. 19. Ratio of total to peak flux as a function of signal to noise of both unresolved (blue) and resolved (red) sources in the combined catalogue.

deviation σ_R ,

$$\sigma_R = \sqrt{\left(\frac{\sigma_S}{S}\right)^2 + \left(\frac{\sigma_{S_{\text{peak}}}}{S_{\text{peak}}}\right)^2}. \quad (14)$$

We take both the uncertainties σ_S and $\sigma_{S_{\text{peak}}}$ as the sum in quadrature of their errors as determined by PYBDSF and a calibration error of 3%. The magnitude of the error in calibration is motivated by assessing Gaussian fits of bright unresolved sources and the flux density offset determined in Sect. 2.5. Using these quantities, the compactness criterion is then

$$\ln\left(\frac{S}{S_{\text{peak}}}\right) > 1.25\sigma_R. \quad (15)$$

The factor of $1.25\sigma_R$ encloses 95% of sources below $S/S_{\text{peak}} = 1$, so with the symmetry of the Gaussian distribution, 95% of all unresolved sources should be correctly identified with this criterion. As can be appreciated in Fig. 19, according to this metric, 50% of all sources in the combined catalogue are resolved.

4.5. Catalogue columns

In the final catalogue, the majority of the columns are preserved from the PYBDSF source catalogues. Additional columns are however added in subsequent steps where required. Our aim is to only add information, and not remove any. This means that, for example, sources can be flagged as artefacts, but will still be present in the catalogue. When performing additional corrections on source flux densities and spectral indices, the correction factors are inserted into the catalogues so that the original values can be easily reproduced. The catalogue has 49 columns in total. Several lines of the catalogue are shown in Table B.1 as an example.

- **Pointing_id** - The ID of the pointing where the source has been found formatted as PT-JHHMM±HHMM.
- **Source_name** - Name of the source, following IAU convention, formatted as JHHMMSS.S±HHMMSS.S with prefix MALS.
- **Source_id** - Source ID as assigned by PYBDSF.
- **Isl_id** - Island ID as assigned by PYBDSF.
- **RA** and **DEC** (and errors) - The J2000 position of the source, defined as the centre of the composite Gaussian of the source, and associated errors.
- **Sep_PC** - Distance of the source from the pointing centre.
- **Total_flux** (and error) - Total flux density of the source and associated error.
- **Flux_correction** - The correction factor for residual primary beam effects on the flux density of the source.
- **Peak_flux** (and error) - Measured peak flux of the source and associated error.
- **Spectral_index** (and error) - Spectral index of the source, measured from the spectral index image, and associated error.
- **Spectral_index_correction** - The correction factor for residual primary beam effects on the spectral index of the source.
- **RA_max** and **DEC_max** (and errors) - Position of maximum intensity of the source and associated errors.
- **Maj**, **Min**, and **PA** (and errors) - FWHM of the major axis, minor axis and position angle of the source fit by PYBDSF and associated errors.

- **DC_Maj**, **DC_Min**, **DC_PA** (and errors) - FWHM of deconvolved major axis, minor axis, and position angle, and associated errors.
- **Isl_Total_flux** (and error) - Total integrated flux of the island in which the source is located, and associated error.
- **Isl_rms** - Average background RMS noise of the island in which the source is located.
- **Isl_mean** - Average background mean value of the island in which the source is located.
- **Resid_Isl_rms** - Average residual background RMS noise of the island in which the source is located.
- **Resid_Isl_mean** - Average residual background mean value of the island in which the source is located.
- **S_Code** - Value generated by PYBDSF indicating whether a source is: fit by a single Gaussian ('S'), fit by multiple Gaussians ('M'), or one of multiple sources on the same island ('C').
- **N_Gaus** - Number of Gaussian components fit to the source.
- **Resolved** - Boolean indicating whether the source is resolved according to the metric defined in Sect. 4.4.
- **Flag_Artifact** - Boolean indicating whether the source is a likely artefact according to the criterion described in Sect. 3.1.
- **RA_mean** and **DEC_mean** - Mean intensity weighted position of all pixels of the island in which the source is located, measured if a source is fit with multiple Gaussians.
- **Cutout_Spectral_index** - Intensity weighted average spectral index of all pixels of the island in which the source is located, measured if a source is fit with multiple Gaussian components.
- **Cutout_Total_flux** - Total flux density of all pixels above the island threshold, measured if source fit with multiple Gaussian components.
- **Cutout_flag** - Flag assigned to cutout in certain conditions: the mean position falls outside the island ('M'), the position of the brightest pixel does not correspond to the maximum position measured by PYBDSF ('C'), the difference between **Cutout_total_flux** and **Isl_Total_flux** is more than 20% ('F').
- **Cutout_class** - Classification assigned at visual inspection as described in Sect. 3.1, indicating whether a source is well described by the Gaussian model ('G'), is better described by the island characteristics ('I'), is better described by a single Gaussian component ('P'), or an artefact ('A').

5. Source characteristics

When considering source counts in the radio regime, extra care must be taken in understanding the population of sources that is being probed. Depending on observing frequency and flux density, different source populations may appear in the sample. Our reference for source counts are the SKADS simulations, as the simulated sample is built up by different source populations. Given the theoretical noise limit of $10 \mu\text{Jy beam}^{-1}$, we can expect to detect sources down to $50 \mu\text{Jy}$. As shown in Wilman et al. (2008), the radio population is dominated by AGN above flux densities of 1 mJy, while below that star-forming galaxies start to make up a significant fraction of the source counts. There are several important distinctions between these source types that can influence source counts and a dipole measurement. Among them is source morphology, as multi-component sources can easily be mistaken for multiple separate sources, biasing number counts. Star-forming galaxies are primarily found at lower




Fig. 20. Fractions of sources fit with varying numbers of Gaussians as a function of flux density. The number of components fit to sources increases steadily towards high flux density, but flattens out around 100 mJy, which can be caused by sources splitting up at these flux densities.

flux densities and can be morphologically described by a single component, such that we expect them to appear as faint point sources in our fields. Consequently, these sources can be easily counted as they are unambiguously unique sources and are thus statistically independent. At higher flux densities however, some sub-classes of AGN, such as Fanaroff-Riley type I (FRI, core-dominated) and type II (FR II, lobe-dominated) sources (Fanaroff & Riley 1974), can boast extended structures that can complicate automated source finding methods.

5.1. Extended sources in PYBDSF

PYBDSF operates by fitting Gaussians to sources, which is an effective method for most radio sources, but breaks down in sources with more complex structure. PYBDSF offers multiple ways of improving the fit to extended sources, as specified in Sect. 3, but this chiefly improves detection of extended and more diffuse sources. To ensure that complex sources are accurately fit by a combination of Gaussians, we employ a special recipe for these types of sources, which make up 8% of the all sources found in the fields. These sources are flagged in our workflow for visual inspection upon which the Gaussian fit is assessed as described in Sect. 3.1.

However, automated source finding algorithms will only recognise objects as a single source if they are closely connected, and will therefore fail on a subset of sources. This effect is strongest for FR II sources, as increased luminosity in the lobes makes them appear as separate radio sources. Source association is one of the outstanding problems in radio astronomy, and beyond the scope of this paper. Instead, we try to characterise this effect and how it might influence source counts. As the components of these sources are not statistically independent, they will naturally bias source counts. To get an estimate of how complex sources are fit by PYBDSF, we look at the number of components fit to sources as a function of flux density in Fig. 20. We see a steady increase in the number of Gaussian components at higher flux densities; however, at ≥ 100 mJy the increase in components flattens out. This may indicate that at these flux densities we are seeing all the emission from these sources, and this is the ‘true’ distribution of components. However, an alternative explanation is that at these flux densities some extended sources no longer have connecting emission and are not recognised as

single sources anymore. In this case the number of components keep increasing, but individual components will split off and be detected as separate sources, effectively keeping the number of components per source the same. Another indication of this happening might be seen in the differential source counts in Fig. 24, where there is an excess in source counts above 100 mJy, relative to the expected values.

To get an alternative measure of this, we look at FR II galaxies in SKADS and the separation of their components. For an upper limit estimate on how many sources we expect to split up, we count components as separated when the distance between them exceeds $6.5''$, which is the minor axis of the average clean beam. With this, 11% of sources in the range 10 mJy–1 Jy have two separated components. Furthermore, 6.3 and 24% of sources in the ranges 10–100 mJy and 100 mJy–10 Jy, respectively, have three separated components. Outside of these ranges the fractions are negligible. Doubling the distance required for separation mostly exchanges the amount of triple component sources for double component sources in the range 100 mJy–1 Jy. The amount of triple component sources does not change in the range 1–10 Jy, indicating that the brightest sources are also the largest and most likely to separate. If we define the excess fraction of sources detected as $f_n = \frac{\tilde{n}-n}{n}$, where \tilde{n} is the amount of sources detected counting separate components and n the actual amount of sources, $f_n = 0.4$ at 10–100 mJy, $f_n = 0.9$ at 100 mJy–1 Jy, and $f_n = 0.7$ at 1–10 Jy. The values of f_n are given for each of the flux density bins used to determine number counts in Table 3. The maximum value f_n can take is 2, when a bin is entirely occupied by sources with three separate components. From this we can conclude that this effect is more important at higher flux densities, and is most significant at flux densities ≥ 100 mJy, which contains only a very tiny subset (102 sources, 0.6%) of the full catalogue.

5.2. Spectral indices

Due to different emission mechanisms and sources of emission, there can be differences in spectral index distribution between star-forming galaxies and AGN. Additionally, the Doppler shift observed as a consequence of the motion of the observer induces a change in observed flux density that depends on the spectral index of the source. Thus, the spectral indices of sources influences the magnitude of the radio dipole. In general, dipole studies assume a single value for spectral index based on the physics of synchrotron emission, $\alpha \approx -0.75$ near 1 GHz (e.g. Rubart & Schwarz 2013; Tiwari & Jain 2013; Siewert et al. 2021). Measuring the spectral index of sources generally requires either large bandwidth or measurements at different frequencies, which in turn requires high S/N to ensure that sources are detected at both ends of the frequency range. With the large bandwidth (802.5 MHz) of MALS we are able to create spectral index images, as described in Sect. 2.3, and measure spectral indices of sources in the catalogue.

To compare the source types, we look at the spectral indices of all sources with respect to flux density. While we cannot completely separate these source types based on flux density or spectral index, these source populations are labelled accordingly in the SKADS simulated sample (Wilman et al. 2008). Figure 21 shows the distribution of spectral indices measured compared to the populations in the SKADS simulated sample. Here, the AGN are separated into FRI (orange) and FR II (red) sub-populations, as are star-forming galaxies separated into ‘starburst’ (light blue) and ‘normal’ (dark blue) galaxies. It is noteworthy that the spectral index distribution of AGN boasts




Fig. 21. Distribution of spectral indices of MALS sources. Left: MALS spectral indices (black) compared to AGN (red) and star-forming galaxies (blue) from SKADS as a function of flux density. Right: MALS spectral indices as a function of flux density, sources with S/N above 20 are coloured by S/N . The median value and error of spectral indices of different flux bins are indicated by red error bars, indicating that at lower flux densities spectral indices tend towards lower values.

two peaks, corresponding to lobes at $\alpha \approx -0.75$, and cores at $\alpha \approx -0.25$. These peaks are respectively dominated by FRII and FRI galaxies, but there is cross-contamination present as we have plotted source components rather than combined source characteristics.

The right plot of Fig. 21 highlights the MALS sources with high S/N , which are more likely to be detected across the full band and thus have more reliable spectral index measurements. Though less tight, the peak at $\alpha \approx -0.75$ is present in the MALS sample. The distribution is broad enough to have mixed with the peak $\alpha \approx -0.25$, so this second peak is possibly lost in the data. To see if we can retrieve these two populations, we take all sources with $S/N > 20$ and assert that the spectral index of the sources above and below $\alpha = -0.5$ represent FRI and FRII sources, respectively. With this definition, there are 626 FRI sources, of which 54% are resolved and 16% are fit with multiple Gaussian components. The brightest of these can be appreciated in Fig. A.1, showing mostly point sources or sources where core emission dominates. FRII sources are more numerous, with 2581 present in the catalogue, of which 75% resolved and 30% fit with multiple Gaussian components. A set of the brightest FRII are also shown in Fig. A.2, showing more sources with two or more components representing radio lobes. These results show that the expected dichotomy in morphology between these sources is indeed present, with FRII sources being more likely identified as extended and/or resolved.

At lower flux densities there appears to be a discrepancy between measured and theoretical spectral indices. We see in Fig. 21 that not only is there a wider distribution in the MALS data, spectral indices are steeper at lower flux densities compared to the SKADS sources. The median spectral index for sources with $S > 1$ mJy is $\alpha = -0.76$, while for sources with $S < 1$ mJy it is $\alpha = -1.17$. Because of the high sensitivity that is needed, spectral indices are not commonly measured at lower flux densities. Looking at deep field surveys however, we see that this result is inconsistent with the spectral indices found in the XMM-LSS/VIDEO deep field (Heywood et al. 2020), where it is found that at lower flux densities spectra flatten out. The $S/N > 20$ sources shown in the right plot of Fig. 21, though

increasing in spread at lower flux densities, are not affected by the same bias.

We further investigate the bias seen in low S/N sources, and verify the corrections made in Sect. 4.1. We compare our spectral indices to those generated by comparing flux densities in SPW2 (1.0 GHz) and SPW9 (1.38 GHz) of the same MALS data by Deka et al. (2022). Smolčić et al. (2017) find a discrepancy between the spectral indices generated by MTMFS deconvolution and those generated by comparing flux densities at different frequencies, so we make the same comparison in Fig. 22, showing the offset between the MTMFS and SPW derived spectral indices. The median offsets for both corrected (blue) and uncorrected (red) are shown, indicating that spectral indices have been properly corrected for residual primary beam effects. Though the offset trends negatively at lower flux densities, it is well within the uncertainties. Overall, the offset between spectral indices is -0.06 ± 0.92 (0.16 ± 0.94 without corrections) for all sources, and -0.01 ± 0.78 for $S/N > 20$ sources, agreeing well between the catalogues. There is no systematic effect seen of the corrected spectral indices with respect to the distance to the pointing centre, indicating no residual primary beam contribution. Deka et al. (2022) observe an overall flattening at low S/N compared to our overall steepening, creating a discrepancy that is clearly showing at lower flux densities. Overall spectral indices appear to be reliable down to mJy flux densities, or S/N of 20. Considering the flux densities of these sources, it is unlikely that many star-forming galaxies are included in the high S/N sample, precluding an analysis of these sources.

5.3. Number counts

Now it is left for us to assert that we have the necessary number counts for a dipole measurement. Extrapolating from the combined catalogue of the ten pointings, a catalogue of the first 391 pointings is expected to carry $\sim 650\,000$ sources, enough to produce a dipole estimate if most sources can be used. However, because our pointings are inhomogeneous, both in terms of internal structure as well as with respect to other pointings, we have to assess to which extent this affects number counts and whether




Fig. 22. Offsets of spectral indices measured from the wideband MTMFS images with respect to spectral indices derived from processing the full bandwidth using 15 individual SPW images from [Deka et al. \(2022\)](#) as a function of flux density. The spectral indices are calculated by using the SPW2 and SPW9 images. Binned median offsets are shown (blue), along with the offsets without the correction applied in Sect. 4.1 (red), showing that the spectral indices are properly corrected.




Fig. 23. Differential source counts from the SKADS simulations. The complete SKADS sample is shown (black), as well as the sample extracted from SKADS and injected into the images (blue). The uncorrected number counts (orange) indicate the sources detected by the source finding routine, which are then corrected with the RMS noise coverage (green).

corrections can be made to homogenise the catalogues. The most common method of comparing number counts to other surveys or simulations is to compute differential source counts, which describes the number of sources dN within a given flux density bin $S + dS$ per steradian on the sky. This is usually multiplied by $S^{5/2}$, which would yield a flat curve in a static Euclidean Universe ([Condon & Ransom 2016](#)).

To verify that we would be able to retrieve the correct number counts, we repeated the experiment carried out in Sect. 3.2, injecting and retrieving sources in the residual images. To simulate a realistic physical distribution of sources, we cut out an area equal to the size of the pointing from the SKADS simulated catalogue. We repeated this experiment for every pointing, each time choosing a random position in the SKADS sample as pointing centre. Figure 23 shows the differential number counts for all the

stages of the experiment. The reference sample from the SKADS simulations (black) represents the full 10×10 degree area simulated in [Wilman et al. \(2008\)](#). Out of the full sample, we cut out ten pointings with the same sky area as the MALS pointings that are injected into the residual images (blue). We performed our source finding routine on these images and saw what number counts we could retrieve. Figure 23 shows that below a few millijansky flux densities, detected source counts begin to fall off (pink), indicating the limit of 100% completeness for the full catalogue. These number counts are normalised by the area coverage of the pointings; however, we can make a simple correction based on the fact that the area coverage is not constant between different flux bins due to varying RMS noise. The actual area covered in a certain flux bin $S + dS$ can be obtained by taking the RMS noise coverage (as shown in Fig. 4) assuming a detection limit of 5σ . This basic correction (green squares) produces correct number counts down to 100–200 μJy , showing that we can account for completeness of the catalogue down to this flux density. Below this, we reach the absolute sensitivity limit of the pointings, as the RMS noise coverage is so low that it produces diverging results. Above 100 mJy, results are more scattered, mainly because of low number counts at these flux densities and the smaller sky coverage of our pointings (35.7 sq. deg.) compared to SKADS (100 sq. deg.).

Having shown that we can reproduce number counts for a large range of flux densities, we measure differential number counts for the combined MALS catalogue. Both corrected and uncorrected differential source counts are tabulated in Table 3 and shown in Fig. 24, where they are compared to the SKADS simulated sample. We apply a completeness correction (purple hexagons) using either unresolved or resolved completeness based on whether the source is classified as such using the criterion from Sect. 4.4. Once again we also correct number counts using the RMS noise coverage (green squares). For comparison, number counts from the full SKADS simulated sample are also shown, both for the central frequency of 1.27 GHz and for the full frequency band, accounting for the fact that for most sources flux density is not equal across the band. The number counts derived from the MeerKAT DEEP2 image and NVSS by [Matthews et al. \(2021\)](#) are also shown (white diamonds). Error bars are computed taking into account Poisson uncertainties as well as source clustering following [Heywood et al. \(2013\)](#). The same data from Fig. 24 is tabulated in Table 3, showing the numerical values of flux bins, raw number counts, and uncorrected and corrected differential number counts. Though it is not taken into consideration in Fig. 24, the number of false detection per bin is included in Table 3. In all cases the number of false detections is smaller than the uncertainty on the number counts, from which we infer that the purity is a small factor compared to completeness of the catalogue.

As for the simulated sample, the corrections appear to hold down to 100–200 μJy , after which we reach the sensitivity limit and solutions diverge. Until solutions diverge, the completeness and RMS noise coverage corrections produce similar results, indicating that a simple sky coverage correction performs well given the ease with which it can be generated. Given the agreement between the completeness and RMS noise coverage corrections, we can safely say that the major contributor to completeness is the inhomogeneous sky coverage of the pointings. This results in the dependence of completeness on the local noise seen for point sources in Fig. 8. With these corrections applied, Fig. 24 shows a wide range of flux densities where the differential number counts deviate from the expected values. At high flux densities (>100 mJy), we see an increase in sources that can




Fig. 24. Differential source counts from MALS, uncorrected (beige circles) and corrected with completeness (purple hexagons) taking into account whether a source is resolved or unresolved. Lastly, source counts are also corrected with RMS noise coverage (green squares), which for the lowest flux bins goes to zero, causing solutions to diverge. This is all compared to the SKADS source counts, both for the central frequency of 1270 MHz (black), and for the full frequency range (grey area), and to the source counts derived from the MeerKAT DEEP2 image and NVSS from [Matthews et al. \(2021\)](#), (white diamonds). The number counts are tabulated in Table 3.

for some part be attributed to the central sources in the images, as well as single sources being classified as multiple sources, as described in Sect. 5.1. However, we see that number counts, except for the range 20–200 mJy, are higher across the board than what we might expect given theoretical predictions. There is some evidence that the SKADS simulations underestimated the number of star-forming galaxies, causing lower counts at low flux densities compared to what is seen in nature (e.g. [Hale et al. 2023](#)). Though this can explain an offset at the lowest flux densities, this effect would only be significant up to millijansky flux densities, whereas our number counts are higher up to an order of magnitude above that. An alternative explanation is that, with its selection of high flux density sources as pointing targets, MALS is probing overdensities, which naturally boosts the number of sources in the pointings. Lastly, such an effect can also be produced by a systematic overestimation of flux densities, which is an option that cannot be ruled out at this stage. We expect that this offset might also be caused by low number statistics, and may disappear once more data are added.

6. Towards the cosmic radio dipole

With a thorough assessment of the quality of the pointings described in this work we have the opportunity to extrapolate our findings to the larger survey of 391 pointings, both in estimating how many pointings will be needed for a dipole estimate, as well as how to effectively homogenise the catalogues. Given these results, there are however some questions and limitations that remain, and these will have to be addressed in later works.

We first estimate the statistical power we can reach with the observed MALS pointings. As we determined in Sect. 2.6, the demerit score is a strong indicator for pointings with high noise and thus low number counts. Consequently, the demerit score

allows us to make predictions about the quality of other pointings in the survey. Of the ten pointings we have investigated, we consider seven of them to be of good quality based on their noise and source count values. Since all MALS pointings are in the footprint of either NVSS or SUMSS, we use these surveys to match sources and calculate the demerit score to predict quality of the images. Based on this principle Fig. 25 shows the distribution of the demerit scores of the first 391 observed MALS pointings. Defining a quality threshold of $d < 15$ mJy based on the ten pointings we have investigated here, we see that 322 of 391 pointings are below this threshold. The seven good quality pointings average ~ 2000 sources per pointing, meaning that 100 such pointings will result in approximately 2×10^5 sources. If we choose our pointings to properly cover the sky along the axis of the dipole, this is the minimum number of sources required for a 3σ measurement of the cosmic radio dipole, assuming an amplitude equal to that of the CMB dipole.

Based on the demerit scores of the ten pointings and those of the first 391 pointings, we also see that our ten pointings well represent the average pointing, and we can use the differential number counts to extrapolate them to the rest of the survey. This allows us to take into account the completeness of the survey so far. As shown in Figs. 23 and 24, we get correct source counts down to 100–200 μ Jy. Naturally, this means that sources below that flux density cannot be included in a dipole estimate, which leaves 13 663 sources in the combined catalogue. However, the corrections are essentially compensating for the missing sources, yielding effectively $\sim 28\,000$ sources down to 200 μ Jy, meaning even fewer than 100 pointings would again suffice for a 3σ dipole measurement.

In both cases we may expect that the first 391 observed pointings will yield around a million sources, which should yield a dipole measurement at a significance level of 6.5σ , assuming

Table 3. Differential source counts of MALS, including corrected counts using RMS coverage, completeness of unresolved sources, and completeness of resolved sources.

S	S_{mean}	N	N_{false}	Sky coverage	$S^{5/2} \frac{dN}{dS}$	Corrected $S^{5/2} \frac{dN}{dS}$	Corrected $S^{5/2} \frac{dN}{dS}$	f_n
(mJy)	(mJy)			(sq. deg.)	(Jy ^{3/2} sr ⁻¹)	RMS coverage (Jy ^{3/2} sr ⁻¹)	Completeness (Jy ^{3/2} sr ⁻¹)	SKADS
0.1–0.13	0.11	850 ± 33	4	2.2	0.377 ± 0.015	6.06 ± 0.24	2.90 ± 0.11	0.0
0.13–0.17	0.15	1254 ± 43	11	5.3	0.816 ± 0.028	5.51 ± 0.19	3.93 ± 0.13	0.0
0.17–0.22	0.19	1575 ± 51	13	8.7	1.50 ± 0.05	6.16 ± 0.20	5.23 ± 0.15	0.0
0.22–0.28	0.25	1708 ± 54	12	12.5	2.39 ± 0.076	6.82 ± 0.22	6.47 ± 0.18	0.0
0.28–0.36	0.32	1558 ± 50	20	16.6	3.20 ± 0.10	6.89 ± 0.22	6.91 ± 0.19	0.0
0.36–0.46	0.41	1393 ± 46	22	20.9	4.21 ± 0.14	7.18 ± 0.24	7.58 ± 0.21	0.0
0.46–0.6	0.53	1257 ± 43	19	25.4	5.57 ± 0.19	7.83 ± 0.27	8.76 ± 0.25	0.0
0.6–0.77	0.69	1053 ± 38	15	23.3	6.85 ± 0.25	10.5 ± 0.4	9.67 ± 0.30	0.0
0.77–1	0.89	881 ± 34	16	31.6	8.41 ± 0.33	9.49 ± 0.37	11.1 ± 0.4	0.0
1–1.3	1.1	769 ± 31	10	30.9	10.8 ± 0.5	12.5 ± 0.5	13.4 ± 0.5	0.0
1.3–1.7	1.5	635 ± 28	7	34.4	13.1 ± 0.6	13.6 ± 0.6	15.5 ± 0.6	0.0
1.7–2.2	1.9	503 ± 24	4	35.0	15.2 ± 0.7	15.5 ± 0.8	17.3 ± 0.8	0.0
2.2–2.8	2.5	411 ± 21	7	35.3	18.2 ± 1.0	18.4 ± 1.0	20.3 ± 1.0	0.0
2.8–3.6	3.2	318 ± 18	2	35.5	20.7 ± 1.2	20.8 ± 1.2	22.5 ± 1.3	0.0
3.6–4.6	4.1	307 ± 18	3	35.6	29.3 ± 1.8	29.4 ± 1.8	31.3 ± 1.8	0.01
4.6–6	5.3	232 ± 15	6	35.7	32.5 ± 2.2	32.5 ± 2.2	34.3 ± 2.3	0.02
6–7.7	6.9	199 ± 14	1	35.7	40.9 ± 3.0	40.9 ± 3.0	42.7 ± 3.0	0.07
7.7–10	8.9	160 ± 13	1	35.7	48.3 ± 3.9	48.3 ± 4.0	49.6 ± 4.0	0.11
10–13	11	137 ± 12	1	35.7	60.7 ± 5.3	60.7 ± 5.3	60.7 ± 5.3	0.16
13–17	15	136 ± 11	0	35.7	88.5 ± 7.8	88.5 ± 7.8	88.5 ± 7.8	0.22
17–22	19	84 ± 9	1	35.7	80.2 ± 8.9	80.2 ± 8.9	80.2 ± 8.9	0.32
22–28	25	100 ± 10	0	35.7	140 ± 14	140 ± 14	140 ± 14	0.5
28–36	32	48 ± 6	0	35.7	98.7 ± 14.4	98.7 ± 14.4	98.7 ± 14.4	0.5
36–46	41	50 ± 7	0	35.7	151 ± 22	151 ± 23	151 ± 22	0.47
46–60	53	36 ± 6	0	35.7	160 ± 27	160 ± 27	160 ± 27	0.63
60–77	69	34 ± 5	0	35.7	221 ± 38	221 ± 38	221 ± 38	0.73
77–100	89	24 ± 4	0	35.7	229 ± 47	229 ± 47	229 ± 47	0.82
100–130	110	21 ± 4	0	35.7	294 ± 65	294 ± 65	294 ± 65	0.93
130–170	150	19 ± 4	0	35.7	391 ± 90	391 ± 90	391 ± 90	1.0
170–220	190	13 ± 3	0	35.7	392 ± 109	392 ± 109	392 ± 109	0.67
220–280	250	10 ± 3	0	35.7	443 ± 140	443 ± 140	443 ± 140	0.75
280–360	320	10 ± 3	0	35.7	650 ± 206	650 ± 206	650 ± 206	0.89
360–460	410	9 ± 3	0	35.7	859 ± 287	859 ± 287	859 ± 287	1.9
460–600	530	6 ± 2	0	35.7	841 ± 344	841 ± 344	841 ± 344	0.75
600–770	690	5 ± 2	0	35.7	1030 ± 460	1030 ± 460	1030 ± 460	0.5
770–1000	890	3 ± 1	0	35.7	906 ± 523	906 ± 523	906 ± 523	1.5
1000–1300	1100	4 ± 2	0	35.7	1770 ± 890	1770 ± 890	1770 ± 890	1.2
1300–1700	1500	1 ± 1	0	35.7	650 ± 650	650 ± 650	650 ± 650	0.75
1700–2200	1900	1 ± 1	0	35.7	955 ± 955	955 ± 955	955 ± 955	0.0

Notes. Counts are normalised for the sky area of 35.7 degrees. Raw counts (N) and number of false detections in (N_{false}) per bin are also given. The excess fraction of sources due to separated components f_n is also given, based on a separation distance of $6.5''$ of FR II sources in SKADS.

an adequate coverage of the dipole axis and a dipole amplitude equal to that of the CMB (Ellis & Baldwin 1984). Though it is clear that MALS will deliver the number counts needed for a significant measurement of the cosmic radio dipole, the larger challenge is making a measurement while accounting for the systematics present in the survey. In order to thus successfully measure the dipole with MALS, the way forwards is to build on the corrections introduced in Sect. 5.3.

6.1. Compiling a homogeneous catalogue

We have effectively shown that we can make a unified description of the properties of the average MALS pointing, which

should now allow us to homogenise the catalogue. Though the structure of the survey, with deep coverage over distinct patches of sky, appears to not lend itself especially well to large-scale cosmology, the fact that these pointings are all equal area by design allows for straightforward discretisation. In the simplest use case, each pointing of MALS can thus be treated as a single unit simply containing N number of sources. When measuring number counts over a full sky, inhomogeneities between the pointings induce higher order multipoles in the data that will spill over into a dipole measurement. To homogenise the data and get an unbiased estimate of the dipole, we must account for the individual differences between the pointings and calculate the corrected effective number counts N_{eff} .




Fig. 25. Demerit scores calculated for the 391 currently observed MAL sources, using bright (>100 mJy) sources retrieved from their corresponding surveys (NVSS and SUMSS). The dotted line indicates a quality threshold of $d < 15$ mJy based on the quality of the pointings inspected in this paper.

To get to effective number counts, the starting point is the corrections to number counts as shown in Fig. 24 that are seen to largely compensate for incomplete catalogues. We can extend this treatment by assigning a ‘completeness factor’ to each individual source. As discussed in Sect. 5.3, the largest contributor is the incomplete sky coverage, which is well modelled with the RMS noise coverage of the pointings. Immediately we can disentangle the completeness into a detection probability, $P(\text{det})$, and sky coverage, Ω ,

$$\text{Comp}(S, Q_A, \sigma) = P(\text{det}|S, Q_A)\Omega(\sigma). \quad (16)$$

Here the detection probability has a power law dependence on source size Q_A (linear in log-log space, right plot of Fig. 16), and assuming Gaussian errors on the flux density, detection probability should follow a Gaussian cumulative distribution function or similar sigmoid function⁷ as a function of flux density. Sky coverage is exclusively determined by the local RMS σ of the source, encoding the noise structure of the pointing. This completeness factor can largely correct for the inhomogeneities present in the catalogues, but we can make it even more robust by including information on other investigated quantities. Using information on purity, we can define a ‘purity factor’,

$$\text{Purity}(S, \rho, \text{pointing}) = \text{Purity}(S, \rho)\text{Purity}_{\text{pointing}}, \quad (17)$$

which indicates how likely a source is to be a true positive. This depends on distance from the pointing centre ρ , flux density S (Fig. 11) and has a multiplicative factor that indicates a purity level that is different per pointing (Fig. 18). Finally, individual sources have associated uncertainties that can be used to weigh each source accordingly. An obvious choice is a weight based flux density, as the uncertainty in flux density σ_S is dependent on flux density S (Fig. 17), which can be combined with a potential flux density scale error ΔS . We are not limited to one uncertainty factor, and a second choice that is relevant for a dipole measurement is the uncertainty in position σ_ϕ , which in absence of any systematics (as we see in Fig. 12) is simply equal to the measured

⁷ Sigmoid is the collective name of functions following ‘S’-shaped curves, which are well suited to describe the detection probability of data near the detection boundary.

uncertainty $\Delta\phi$. Combining all these measures, we can assign a weight to sources based on the quantities laid out,

$$w_S = \sigma_S^{-1}(S, \Delta S), \quad (18)$$

$$w_\phi = \sigma_\phi^{-1}(\Delta\phi), \quad (19)$$

$$w_{\text{eff}} = w_\phi w_S \frac{\text{Purity}(S, \rho, \text{pointing})}{\text{Comp}(S, Q_A, \sigma)}. \quad (20)$$

$$= w_\phi w_S \frac{\text{Purity}(S, \rho)\text{Purity}_{\text{pointing}}}{P(\text{det}|S, Q_A)\Omega(\sigma)}. \quad (21)$$

The effective weight factor, w_{eff} , fulfils a dual purpose in estimating the number counts. The completeness and purity factors correct the number counts, while the weights from the flux density and position errors then serve as a quality measure for each source, allowing us to measure the effective number density of the individual pointings,

$$N_{\text{eff}} = \frac{\sum_i^n w_{\text{eff},i}}{\sum_i^n w_{\phi,i}w_{S,i}}. \quad (22)$$

6.2. Limitations and future prospects

With this prescription, the systematic effects that we have characterised can be accounted for when computing the number counts and estimating the dipole. There remain however some effects that have not been explicitly characterised that could influence a dipole measurement. By checking the corrections to number counts as we did in Sect. 5.3 on simulated data, we essentially calibrated the corrections on the SKADS sample, which as a simulation might not perfectly represent the number counts found in nature (this can be plainly seen in Fig. 24, where SKADS number counts do not always agree with the counts from Matthews et al. 2021). While this can introduce an unknown error into the process, the error is expected to be in overall number counts, and therefore not directionally dependent. Similarly, an important aspect of MAL is the selection of the pointings, as every pointing has a bright radio source at the centre. Although pointings are distributed isotropically, this is not equivalent to a random selection as bright central sources are more likely to be embedded in overdensities. This effect seems to be very pronounced in our measured source counts already, which are larger than expected. Again, this effect is expected to be directionally independent, but whether this is truly the case remains to be determined. Finally, the depth of MAL might be to its detriment when measuring a dipole, as the reached depth of 200 μJy probes into the population of starburst and normal galaxies. The brightest sub-population of these fainter sources occurs at lowest redshifts, which exhibits stronger clustering than AGN. To what extent this affects a dipole measurement is explored in Bengaly et al. (2019), who perform several redshift cuts and a significant improvement is made with $z_{\text{cut}} = 0.1$ compared to including all sources. This is a more stringent cut than Blake & Wall (2002), who claim to eliminate local clustering effects with $z_{\text{cut}} = 0.03$. Although no direct redshift information is available from the MAL data, we will be able to investigate the effect of clustering due to nearby star-forming galaxies using photometric redshifts from current surveys such as the Dark Energy Camera Legacy Survey (DeCALs; Blum et al. 2016), and the *Rubin* Observatory Legacy Survey of Space and Time (LSST; LSST Science Collaboration 2009) in the near future.

It is worth noting that an increase in number counts as seen in Fig. 24 can also be (partially) caused by a systematic flux density offset. The results seen in Fig. 13 hint to a systematic flux density offset with respect to NVSS, which is worth investigating with the larger MALS catalogue. If this offset turns out to be indeed significant, there are many possible explanations given that such a systematic effect could potentially be introduced at many points in the data processing pipeline. We have already shown in Sect. 4.3.2 that the island flux density from PYBDSF properly recovers the flux density from simulated sources, which verifies that the source finding step is not inducing a systematic flux density offset. To further narrow down the options, we cross-check our results with those from [Deka et al. \(2022\)](#), which use the same data, calibration pipeline and source finding strategy but show an overall agreement with NVSS. The most notable difference between these catalogues is the fact that we utilise the full band while [Deka et al. \(2022\)](#) only use individual SPWs with 50 MHz bandwidth. Such a difference between results from the full bandwidth and individual SPWs could point to systematic effects introduced in the imaging stage, as we model the emission with two Taylor terms in frequency as opposed to a single Taylor term in case of individual SPW images. Though the ten pointings explored here have already provided a wealth of statistics and insight into systematic effects, these will be further investigated with the full suit of MALS pointings.

With regards to a dipole measurement, there are still some questions that remain to be answered. Throughout this work we have used the estimate by [Ellis & Baldwin \(1984\)](#) of 2×10^5 sources properly distributed along the axis of the dipole. Though the MALS pointings properly cover the axis of the CMB dipole, if the direction of the radio dipole deviates from this, for example towards the northern hemisphere, the coverage of MALS pointings might not be adequate. Furthermore, the exact amount of sources needed for a dipole estimate can vary depending on this coverage and has been differently estimated in different works. [Crawford \(2009\)](#) claims 2×10^6 sources are necessary for a 3σ dipole estimate, which is an order of magnitude more than the number from [Ellis & Baldwin \(1984\)](#). Dipole studies using NVSS have generally reached 3σ significance with 3×10^5 sources (e.g. [Singal 2011](#); [Rubart & Schwarz 2013](#); [Secrest et al. 2022](#)). This is in closer agreement to the [Ellis & Baldwin \(1984\)](#) numbers; however, this significance is only reached because of the anomalously high dipole amplitude. Therefore, the significance with which the dipole can be measured ultimately depends on many factors, including sky coverage, number counts, dipole amplitude, frequency, and the employed estimator. Consequently, another important step towards measuring the cosmic radio dipole with MALS is defining an appropriate dipole estimator (see e.g. [Siewert et al. 2021](#)), which beyond the scope of this paper but will be explored in a future work.

7. Summary and conclusion

In this work we have presented a thorough analysis of the first ten deep continuum pointings of MALS ([Gupta et al. 2016](#)) and have compiled a catalogue with 16 307 sources covering 35.7 square degrees of deep radio sky. We set out to extensively analyse the properties of the first ten pointings of MALS, with the ultimate goal of measuring the cosmic radio dipole. To achieve a measurement of number counts unbiased by the inhomogeneities present between the MALS pointings,

we characterised systematic effects that can influence such a measurement. This assessment of systematic effects in the ten pointings as presented in this work shows that these effects are for the most part predictable and can be properly accounted for. This will eventually not only benefit a dipole measurement, but all continuum science carried out with MALS. In the current literature on the cosmic dipole, there are many examples of systematic effects that limit the sensitivity of these estimates and could not be pushed further due to a lack of information on the inner workings of the surveys that were used. For MALS, we have a complete assessment of the inner workings of the survey, with insight and access into the processing pipeline. Looking forwards, we determine that 100 MALS pointings suffice for a dipole measurement. Paired with the analysis on source characteristics and counts in the pointings, we are poised to perform the most complete dipole estimate of a radio survey thus far.

Calibration and imaging of all the data was carried out through ARTIP. After imaging, we separately created spectral index images and performed primary beam correction averaged over the frequency range on these and the continuum images with a primary beam model derived from holographic measurements. We made an initial assessment of the calibration by checking the flux density scale of the calibrators and central sources, and find that flux densities are consistent with those reported in the literature. We investigated the quality of the images by looking at the RMS noise maps created by PYBDSF. Measuring RMS noise coverage shows that all pointings have similar noise structure, but overall noise levels are offset between the pointings. We quantified this offset with σ_{20} , which gives the noise level at 20% RMS noise coverage for a pointing. To try to explain the difference in pointing quality, we calculated demerit scores for each pointing to estimate the contributions of bright sources to the noise. Though there is a correlation between the noise in a pointing and the demerit score, other factors play a part in introducing a scatter in this relation, especially for pointings with lower demerit scores. As σ_{20} directly describes the overall noise in the image, it is the quality measure of choice for the pointings.

We performed source extraction on all the images using PYBDSF and converted PYBDSF catalogues to full Stokes I catalogues, extending them to include spectral indices and flagging artefacts. We considered sources fit with multiple Gaussian components to be potentially complex and visually inspected them to investigate how well the Gaussian fit describes these sources. We further assessed the quality of the individual pointings and how it affects source finding by measuring completeness, flux recovery, and purity. For completeness and flux recovery, we need to know the intrinsic properties of the sources in the images, so we created mock catalogues of sources from the SKADS simulated sample and injected these into the residual images of the pointings. We assessed completeness for unresolved and resolved sources separately. Using unresolved sources, we assessed completeness as a function of distance from the pointing centre, and for resolved sources we investigated completeness as a function of source size.

Combining the catalogues of the individual pointings, we corrected for residual primary beam effects in the flux densities and spectral indices sources originating from the frequency dependence of the primary beam. To check these corrections and other potential systematic effects, we cross-matched the catalogues with NVSS to check if positions and flux densities were consistent. There is no appreciable astrometric offset; there is an 18% offset in flux density compared to NVSS, but this is still

within the uncertainties. Using σ_{20} as a normalisation factor, we combined completeness measures from the individual pointings and find unified completeness relations that hold for all pointings. Combining flux recovery statistics from all pointings, we find that a systematic bias is present in the integrated flux densities from the fitted Gaussians in the catalogue. This bias is not present in the integrated flux densities of the islands that the sources occupy, making this quantity the logical choice of flux density for the analysis presented here. Combining purity from all pointings, we find that we can account for 20% of false detections with a suitable artefact identification scheme. The remaining false detections make up an increasing fraction of sources farther out from the pointing centre.

As the full catalogues are expected to be populated by various sources types, we assessed how this can influence source counts. Looking at the number of Gaussian components fit to sources, we see that the number of components needed to describe sources increases as a function of flux density until it stagnates at around 100 mJy. This implies that, around this flux, density sources separate and can be counted multiple times; however, as only a tiny percentage of sources are present at these flux densities, it is unlikely to bias the source counts. To further differentiate between source populations, we looked at the spectral indices of sources and find that we can reasonably separate core-dominated FRI from lobe-dominated FR II sources by making a cut at $\alpha = -0.5$. At low flux densities we find that spectral indices are much steeper than expected, which is likely caused by low S/N and inadequate modelling by the MTMFS deconvolution scheme.

Finally, to show that we can account for the systematic effects present in the catalogues, we calculated and corrected differential number counts for a set of simulated catalogues using the SKADS simulated sample. We find that a simple correction using only the RMS noise coverage produces correct number counts down to 100–200 μ Jy. We then computed differential number counts for the full catalogue and corrected number counts using the combined completeness measures found for unresolved and resolved sources, as well as the RMS noise coverage. Once again, corrections seem to hold down to 100–200 μ Jy. Comparing these number counts to the expected number counts from the SKADS sample, we see that our number counts are higher in the full range of probed flux densities. This is independent of corrections, so a likely explanation is that MALS is probing overdense regions, as it targets bright sources. The same effect can, however, be (partially) produced by a systematic flux density offset, which should be taken into consideration given the results on the flux density scale.

Using both the demerit score to predict the quality of other MALS pointings and the corrected number counts of this sample, we show that we will require 100 MALS pointings to reach the necessary number counts for a 3σ measurement of the dipole. Going further, we assert that we can assess the dipole on the level of individual sources using the information on flux density scale, completeness, purity, flux errors, and position errors. The precise implementation is left to later works, along with an exploration of viable dipole estimators.

Acknowledgements. We thank the anonymous referee for their useful comments and feedback. We thank Charles Walker for his insightful comments and feedback on the text. J.D.W. acknowledges the support from the International Max Planck Research School (IMPRS) for Astronomy and Astrophysics at the Universities of Bonn and Cologne. E.B. acknowledges support by NASA under award number 80GSFC21M0002. P.K. is partially supported by the BMBF project 05A17PC2 for D-MeerKAT II. The MeerKAT telescope is

operated by the South African Radio Astronomy Observatory, which is a facility of the National Research Foundation, an agency of the Department of Science and Innovation. The MeerKAT data were processed using the MALS computing facility at IUCAA (<https://mals.iucaa.in/releases>). The Common Astronomy Software Applications (CASA) package is developed by an international consortium of scientists based at the National Radio Astronomical Observatory (NRAO), the European Southern Observatory (ESO), the National Astronomical Observatory of Japan (NAOJ), the Academia Sinica Institute of Astronomy and Astrophysics (ASIAA), the CSIRO division for Astronomy and Space Science (CASS), and the Netherlands Institute for Radio Astronomy (ASTRON) under the guidance of NRAO. The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

References

- Baars, J. W. M., Genzel, R., Pauliny-Toth, I. I. K., & Witzel, A. 1977, *A&A*, **61**, 99
- Bengaly, C. A. P., Maartens, R., & Santos, M. G. 2018, *J. Cosmology Astropart. Phys.*, **2018**, 031
- Bengaly, C. A. P., Siewert, T. M., Schwarz, D. J., & Maartens, R. 2019, *MNRAS*, **486**, 1350
- Bhatnagar, S., Cornwell, T. J., Golap, K., & Uson, J. M. 2008, *A&A*, **487**, 419
- Blake, C., & Wall, J. 2002, *Nature*, **416**, 150
- Blum, R. D., Burleigh, K., Dey, A., et al. 2016, *AAS Meeting Abstracts*, **228**, 317.01
- Bonaldi, A., Bonato, M., Galluzzi, V., et al. 2019, *MNRAS*, **482**, 2
- Bridle, A. H., & Schwab, F. R. 1999, *ASP Conf. Ser.*, **180**, 371
- Briggs, D. S. 1995, PhD thesis, New Mexico Institute of Mining and Technology, USA
- Colin, J., Mohayaee, R., Rameez, M., & Sarkar, S. 2017, *MNRAS*, **471**, 1045
- Combes, F., Gupta, N., Muller, S., et al. 2021, *A&A*, **648**, A116
- Condon, J. J., & Ransom, S. M. 2016, *Essential Radio Astronomy* (Princeton: Princeton University Press)
- Condon, J. J., Cotton, W. D., Greisen, E. W., et al. 1998, *AJ*, **115**, 1693
- Cornwell, T., & Fomalont, E. B. 1989, *ASP Conf. Ser.*, **6**, 185
- Cornwell, T. J., Golap, K., & Bhatnagar, S. 2005, *ASP Conf. Ser.*, **347**, 86
- Crawford, F. 2009, *ApJ*, **692**, 887
- Darling, J. 2022, *ApJ*, **931**, L14
- de Villiers, M. S., & Cotton, W. D. 2022, *AJ*, **163**, 135
- Deka, P., Gupta, N., & Jagannathan, P. 2022, *ApJ*, submitted
- Eddington, A. S. 1913, *MNRAS*, **73**, 359
- Ellis, G. F. R., & Baldwin, J. E. 1984, *MNRAS*, **206**, 377
- Fanaroff, B. L., & Riley, J. M. 1974, *MNRAS*, **167**, 31P
- Franzen, T. M. O., Banfield, J. K., Hales, C. A., et al. 2015, *MNRAS*, **453**, 4020
- Grobler, T. L., Nunhokee, C. D., Smirnov, O. M., van Zyl, A. J., & de Bruyn, A. G. 2014, *MNRAS*, **439**, 4030
- Gupta, N., Srianand, R., Baan, W., et al. 2016, in *MeerKAT Science: On the Pathway to the SKA*, 14
- Gupta, N., Jagannathan, P., Srianand, R., et al. 2021, *ApJ*, **907**, 11
- Gupta, N., Shukla, G., Srianand, R., et al. 2022, *ApJ*, **929**, 108
- Hale, C. L., Robotham, A. S. G., Davies, L. J. M., et al. 2019, *MNRAS*, **487**, 3971
- Hale, C. L., McConnell, D., Thomson, A. J. M., et al. 2021, *PASA*, **38**, e058
- Hale, C. L., Whittam, I. H., Jarvis, M. J., et al. 2023, *MNRAS*, **520**, 2668
- Heywood, I., Jarvis, M. J., & Condon, J. J. 2013, *MNRAS*, **432**, 2625
- Heywood, I., Hale, C. L., Jarvis, M. J., et al. 2020, *MNRAS*, **496**, 3469
- Holschneider, M., Kronland-Martinet, R., Morlet, J., & Tchamitchian, P. 1989, in *Wavelets*, eds. J.-M. Combes, A. Grossmann, & P. Tchamitchian (Berlin, Heidelberg: Springer Berlin Heidelberg), 286
- Intema, H. T., Jagannathan, P., Mooley, K. P., & Frail, D. A. 2017, *A&A*, **598**, A78
- Lacy, M., Baum, S. A., Chandler, C. J., et al. 2020, *PASP*, **132**, 035001
- Lineweaver, C. H. 1997, in *Microwave Background Anisotropies* (Berlin: Springer), 16, 69
- Longair, M. S., & Scheuer, P. A. G. 1966, *MNRAS*, **133**, 421
- LSST Science Collaboration (Abell, P. A., et al.) 2009, ArXiv e-prints [arXiv:0912.0201]
- Matthews, A. M., Condon, J. J., Cotton, W. D., & Mauch, T. 2021, *ApJ*, **909**, 193
- Mauch, T., Murphy, T., Buttery, H. J., et al. 2003, *MNRAS*, **342**, 1117
- Mauch, T., Cotton, W. D., Condon, J. J., et al. 2020, *ApJ*, **888**, 61

- McConnell, D., Hale, C. L., Lenc, E., et al. 2020, *PASA*, 37, e048
- Mohan, N., & Rafferty, D. 2015, *Astrophysics Source Code Library* [ascl:1502.007]
- Penzias, A. A., & Wilson, R. W. 1965, *ApJ*, 142, 419
- Perley, R. A., & Butler, B. J. 2017, *ApJS*, 230, 7
- Rau, U., & Cornwell, T. J. 2011, *A&A*, 532, A71
- Rengelink, R. B., Tang, Y., de Bruyn, A. G., et al. 1997, *A&AS*, 124, 259
- Rubart, M., & Schwarz, D. J. 2013, *A&A*, 555, A117
- Ryle, M., & Scheuer, P. a. G. 1955, *Proc. R. Soc. London Ser. A Math. Phys. Sci.*, 230, 448
- Secrest, N. J., von Hausegger, S., Rameez, M., et al. 2021, *ApJ*, 908, L51
- Secrest, N. J., von Hausegger, S., Rameez, M., Mohayaee, R., & Sarkar, S. 2022, *ApJ*, 937, L31
- Shimwell, T. W., Tasse, C., Hardcastle, M. J., et al. 2019, *A&A*, 622, A1
- Siewert, T. M., Schmidt-Rubart, M., & Schwarz, D. J. 2021, *A&A*, 653, A9
- Singal, A. K. 2011, *ApJ*, 742, L23
- Singal, A. K. 2019, *Phys. Rev. D*, 100, 063501
- Singal, A. K. 2021, *Universe*, 7, 107
- Smolčić, V., Novak, M., Bondi, M., et al. 2017, *A&A*, 602, A1
- Taylor, A., & Legodi, L. 2021, Data Products for Polarization of MeerKAT Calibrators, 1.0 (South African Radio Astronomy Observatory) <https://doi.org/10.48479/an2v-hy97>
- The CASA Team, Bean, B., Bhatnagar, S., et al. 2022, *PASP*, 134, 114501
- Tiwari, P., & Jain, P. 2013, *MNRAS*, 447, 2658
- Tiwari, P., Kothari, R., Naskar, A., Nadkarni-Ghosh, S., & Jain, P. 2015, *Astropart. Phys.*, 61, 1
- Wilman, R. J., Miller, L., Jarvis, M. J., et al. 2008, *MNRAS*, 388, 1335

Appendix A: Cutouts of bright FRI and FRII sources




Fig. A.1: Selection of brightest sources in the combined catalogue with $\alpha \geq -0.5$. As discussed in Sect. 5.2, this range of spectral indices is expected to come from emission originating in AGN cores, which are generally point sources. The majority of this sample indeed is unresolved or can be seen to dominate the emission from their associated lobes.




Fig. A.2: Selection of brightest sources from the combined catalogue with $\alpha < -0.5$. As discussed in Sect. 5.2, this range of spectral indices is generally associated with synchrotron, which is the dominant emission mechanism in radio lobes. Though a number of sources in this sample are unresolved, many show the two component structure characteristic of FRIIs or even more complex extended structure.

Appendix B: Table of sources

Table B.1: Example of the final source catalogue structure. Columns are as described in Section 4.5.

Pointing_id	Source_name	Source_id	Isl_id	RA J2000, °	E_RA °	DEC J2000, °	E_DEC °	Sep_PC °
PT-J0001-1540	J000012.63-154312.0	1607	1618	0.052609	6.82e-05	-15.720001	9.42e-05	0.36
PT-J0006+1728	J000409.33+163842.5	1285	1331	1.038859	8.34e-05	16.645139	1.36e-04	1.04
PT-J0126+1420	J012317.38+133309.1	1504	1530	20.822411	1.70e-04	13.552529	1.65e-04	1.06
PT-J0240+0957	J023823.97+094637.3	888	904	39.599861	3.05e-05	9.777035	5.93e-05	0.54
PT-J0249+0440	J024648.18+035955.5	1469	1483	41.700764	8.47e-05	3.998738	1.61e-04	0.98
PT-J0249-0759	J024633.91-073732.5	2518	2529	41.641290	1.64e-04	-7.625700	2.18e-04	0.83
PT-J1133+0015	J113059.76+000708.0	708	717	172.749000	2.06e-04	0.118892	2.53e-04	0.53
PT-J1232-0224	J122945.15-021418.1	507	517	187.438138	1.29e-04	-2.238372	1.77e-04	0.59
PT-J1312-2026	J130900.75-204924.6	2326	2334	197.253139	4.04e-05	-20.823501	5.03e-05	0.82
PT-J2023-3655	J202008.96-371033.3	2055	2067	305.037343	2.38e-04	-37.175915	1.73e-04	0.77

Total_flux mJy	E_Total_flux mJy	Peak_flux mJy beam ⁻¹	E_Peak_flux mJy beam ⁻¹	Spectral_index	Spectral_index_correction	E_Spectral_index
0.409	0.050	0.273	0.022	-1.90	0.10	0.21
4.564	0.458	2.101	0.151	-0.10	0.80	0.09
19.007	1.561	2.288	0.121	-0.53	0.83	0.07
4.692	0.368	2.575	0.062	-0.74	0.21	0.05
1.418	0.231	0.985	0.102	0.46	0.72	0.16
0.421	0.113	0.285	0.049	-0.94	0.51	0.27
0.423	0.114	0.254	0.046	-1.51	0.21	0.10
0.530	0.121	0.394	0.057	-2.36	0.25	0.10
1.761	0.117	1.136	0.050	0.09	0.50	0.07
0.511	0.116	0.306	0.047	-1.56	0.43	0.23

RA_max J2000, °	E_RA_max °	DEC_max J2000, °	E_DEC_max °	Maj "	E_Maj "	Min "	E_Min "	PA °	E_PA °
0.052609	6.82e-05	-15.720001	9.42e-05	9.35	0.80	7.78	0.57	6.66	19.55
1.038859	8.34e-05	16.645139	1.36e-04	14.55	1.16	10.66	0.70	171.83	11.22
20.824116	1.70e-04	13.551953	1.65e-04	19.38	1.50	17.70	1.33	38.22	36.76
39.601394	3.05e-05	9.776665	5.93e-05	14.21	0.51	9.25	0.25	80.33	3.76
41.700764	8.47e-05	3.998738	1.61e-04	11.36	1.38	7.74	0.69	10.10	13.35
41.641290	1.64e-04	-7.625700	2.18e-04	10.19	1.97	7.82	1.22	154.03	30.80
172.749000	2.06e-04	0.118892	2.53e-04	11.83	2.42	8.42	1.34	146.13	26.09
187.438138	1.29e-04	-2.238372	1.77e-04	9.72	1.53	7.90	1.04	162.61	31.48
197.253139	4.04e-05	-20.823501	5.03e-05	9.35	0.43	8.03	0.33	17.92	12.99
305.037343	2.38e-04	-37.175915	1.73e-04	12.52	2.08	9.91	1.38	72.24	30.76

DC_Maj "	E_DC_Maj "	DC_Min "	E_DC_Min "	DC_PA °	E_DC_PA °
5.60	0.80	4.16	0.57	31.11	19.55
9.03	1.16	8.62	0.70	161.03	11.22
17.71	1.50	15.05	1.33	64.76	36.76
0.00	0.51	0.00	0.25	0.00	3.76
6.93	1.38	3.53	0.69	22.80	13.35
6.49	1.97	3.62	1.22	141.44	30.80
8.15	2.42	4.45	1.34	133.16	26.09
4.97	1.53	3.85	1.04	135.68	31.48
6.02	0.43	4.12	0.33	44.66	12.99
8.28	2.08	4.65	1.38	29.77	30.76

Table B.1: continued

Isl_Total_flux mJy beam ⁻¹	E_Isl_Total_flux mJy beam ⁻¹	Isl_rms mJy beam ⁻¹	Isl_mean mJy beam ⁻¹	Resid_Isl_rms	Resid_Isl_mean	S_Code	N_Gaus
0.369	0.034	0.020	-0.004	0.007	-0.004	S	1
4.343	0.276	0.141	-0.021	0.121	-0.002	S	1
11.940	0.416	0.121	-0.010	0.121	-0.010	M	2
3.954	0.160	0.062	0.005	0.021	0.004	M	2
1.231	0.136	0.098	-0.020	0.019	-0.019	S	1
0.349	0.060	0.047	-0.000	0.009	-0.000	S	1
0.384	0.064	0.043	0.003	0.003	0.003	S	1
0.426	0.065	0.054	-0.001	0.006	-0.001	S	1
1.613	0.082	0.046	-0.019	0.035	-0.014	S	1
0.373	0.056	0.043	-0.008	0.007	-0.008	S	1

Resolved	Flag_artifact	RA_mean J2000, °	DEC_mean J2000, °	Cutout_Spectral_index	Cutout_Total_flux mJy	Cutout_flag	Cutout_class
True	False	–	–	–	–	–	–
True	False	–	–	–	–	–	–
True	False	20.822423	13.552525	10.482	-0.24	C	I
True	False	39.599830	9.777032	3.583	-0.75	C	P
True	False	–	–	–	–	–	–
False	False	–	–	–	–	–	–
True	False	–	–	–	–	–	–
False	False	–	–	–	–	–	–
True	False	–	–	–	–	–	–
False	False	–	–	–	–	–	–

Paper II

The paper *Wagenveld et al. (2023b)*, *A&A*, 675, A72, 14 was published in *Astronomy & Astrophysics* under the Creative Commons CC-BY 4.0 license, and is reproduced in its original form below.

The cosmic radio dipole: Bayesian estimators on new and old radio surveys

J. D. Wagenveld¹, H.-R. Klöckner¹, and D. J. Schwarz²

¹ Max-Planck Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany
e-mail: wagenveld@mpi-fr-bonn.mpg.de

² Fakultät für Physik, Universität Bielefeld, Postfach 100131, 33501 Bielefeld, Germany

Received 21 February 2023 / Accepted 24 May 2023

ABSTRACT

The cosmic radio dipole is an anisotropy in the number counts of radio sources and is analogous to the dipole seen in the cosmic microwave background (CMB). Measurements of source counts of large radio surveys have shown that, although the radio dipole is generally consistent in direction with the CMB dipole, the amplitudes are in tension. These observations present an intriguing puzzle, namely the cause of this discrepancy, with a true anisotropy breaking with the assumptions of the cosmological principle, invalidating the most common cosmological models that are built on these assumptions. We present a novel set of Bayesian estimators to determine the cosmic radio dipole and compare the results with those of commonly used methods applied to the Rapid ASKAP Continuum Survey (RACS) and the NRAO VLA Sky Survey (NVSS) radio surveys. In addition, we adapt the Bayesian estimators to take into account systematic effects known to influence large radio surveys of this kind, folding information such as the local noise floor or array configuration directly into the parameter estimation. The enhancement of these estimators allows us to greatly increase the number of sources used in the parameter estimation, yielding tighter constraints on the cosmic radio dipole estimation than previously achieved with NVSS and RACS. We extend the estimators further to work on multiple catalogues simultaneously, leading to a combined parameter estimation using both NVSS and RACS. The result is a dipole estimate that perfectly aligns with the CMB dipole in terms of direction but with an amplitude that is three times as large, and a significance of 4.8σ . This new dipole measurement is made to an unprecedented level of precision for radio sources, which is only matched by recent results using infrared quasars.

Key words. large-scale structure of Universe – galaxies: statistics – radio continuum: galaxies

1. Introduction

The cosmic microwave background (CMB) and the structures seen therein reveal information about the large-scale structure in the Universe. In addition to the well-studied low-level ($\Delta T/T \sim 10^{-5}$) anisotropies, there is a larger anisotropy seen in the CMB known as the cosmic dipole ($\Delta T/T \sim 10^{-3}$), an effect attributed to the movement of the Solar System with respect to the CMB restframe. The velocity of the Solar System as derived from the amplitude of the CMB dipole is $v = 369.82 \pm 0.11 \text{ km s}^{-1}$ (Aghanim et al. 2020), assuming that the dipole is entirely caused by the motion of the observer with respect to the CMB.

Ellis & Baldwin (1984) first proposed a way to measure the dipole in the number counts of radio sources, as the radio population outside of the Galactic plane mostly consists of extragalactic sources that are expected to be part of and to trace large-scale structure. The first significant measurement of the cosmic radio dipole¹ was reported by Blake & Wall (2002) using the National Radio Astronomy Observatory (NRAO) Very Large Array (VLA) Sky Survey (NVSS, Condon et al. 1998); their result agreed with the CMB dipole within uncertainties. However, subsequent studies using the NVSS found the amplitude of the cosmic radio dipole to be significantly larger than that of the CMB dipole. Singal (2011) found a radio dipole to be

four times larger, which was corroborated by Rubart & Schwarz (2013), and both measurements were obtained with a 3σ significance. In addition to the NVSS measurement, the cosmic radio dipole was measured with other radio surveys, such as the Westerbork Northern Sky Survey (WENSS, Rengelink et al. 1997), the Tata Institute for Fundamental Research (TIFR) Giant Metrewave Radio Telescope (GMRT) Sky Surveys first alternative data release (TGSS ADR1, Intema et al. 2017), and the Sydney University Molonglo Sky Survey (SUMSS, Mauch et al. 2003). Though depending on the survey and employed estimator, the amplitude of the radio dipole is consistently larger than the amplitude of the CMB dipole (see Siewert et al. 2021, for an overview), while the direction of the radio dipole remains consistent with that of the CMB, albeit with considerable uncertainty. It has been argued that failing to take into account the evolution of the spectral index, the magnification bias, or the luminosity function of the source population over cosmic time can potentially bias the radio dipole (Dalang & Bonvin 2022; Guandalin et al. 2022), although so far this evolution has not been observed (e.g. Böhme et al. 2023).

Most significant are the results from Secrest et al. (2021, 2022), which find a dipole amplitude over twice that of the CMB at a significance of 4.9σ using Wide-field Infrared Survey Explorer (WISE, Wright et al. 2010) measurements of quasars. Secrest et al. (2022) perform a joint analysis of NVSS radio galaxies and WISE quasars, with a resulting significance of 5.1σ . A Bayesian estimator based on Poisson statistics was recently utilised by Dam et al. (2022) for the first time to

¹ Both ‘cosmic radio dipole’ and ‘radio dipole’ are used in this paper to refer to the dipole observed in the number counts of radio sources.

measure the cosmic radio dipole with the WISE quasar sample from [Secrest et al. \(2021\)](#), yielding a dipole amplitude 2.7 times larger than the CMB dipole with a significance of 5.7σ . Ultimately, a discrepancy between a dipole in the CMB and other observables points to an unknown effect on the data, which is increasingly unlikely to be systematic among all the different probes, as the NVSS and WISE samples for example are independent of one another, both in terms of source population and systematic effects. This points to an unexpected anisotropy in the large-scale structure of the Universe, something which breaks with the core assumptions of the cosmological principle. If genuine, this anisotropy poses a major problem for cosmologies that are based on the Friedmann-Lemaître-Robertson-Walker metric, such as Λ -CDM.

One problem that remains persistent, even with ever larger datasets, is a lack of homogeneity caused by systematic effects that influence the data. Systematic effects that are unaccounted for can greatly bias dipole estimates, and so conservative cuts in the data must be made to eliminate biases as much as possible. Already in the first measurement of the cosmic radio dipole using the NVSS ([Blake & Wall 2002](#)), a persistent systematic effect was identified that causes large differences in source density as a function of declination, inducing an artificial north-south anisotropy in the data. To eliminate this effect and avoid biasing dipole estimates, conservative cuts in flux density have to be made in such catalogues, which greatly reduces the number of usable sources.

[Wagenveld et al. \(2023\)](#) presented a deep analysis of ten pointings from the MeerKAT Absorption Line Survey ([Gupta et al. 2016](#)), with a focus on mitigating biases that could affect a measurement of the cosmic radio dipole. Being able to account for systematic effects allows less strict flux density cuts to be made, increasing the homogeneity of the catalogues and the number of sources that can be used for a dipole estimate. This approach was made possible by having direct access to meta data and data products within the processing steps of the survey from calibration to imaging and source finding, which is commonly not the case in dipole studies. In the present work, we approach this problem from the outside, and show how dipole estimates can be improved with the information present in modern radio catalogues. We present new dipole estimators, constructing likelihoods for estimating dipole parameters that can take this information into account, as well as an estimator that combines different catalogues for an improved dipole estimate. Given the proper information, these estimators are able to account for systematic effects on number counts in radio surveys and remove the need to cut large amounts of data.

This paper is organised as follows. In Sect. 2, we describe the statistics of radio source counts and the dipole effect. In Sect. 3, we introduce the estimators that will be used to infer the radio dipole parameters using the data sets described in Sect. 4. The results obtained are given in Sect. 5. The implications of our results are discussed in Sect. 6, along with some caveats. In Sect. 7, we summarise the findings of our study.

2. Radio source counts and the dipole

The majority of bright sources at radio wavelengths outside of the Galactic plane are active galactic nuclei (AGN) and have a redshift distribution that peaks at $z \sim 0.8$ (e.g. [Condon & Ransom 2016](#)). As such, radio sources are expected to trace the background, and should therefore comply with isotropy and homogeneity on the largest scales. Following the cosmological principle, the surface density of sources should

therefore be independent of location on the sky. The naive expectation is a distribution of radio sources that are independent, identical, and point-like, which defines a Poisson point process. By discretising the sky into regions of finite size, the number of sources per region will follow a Poisson distribution. The probability density distribution

$$p(n) = \frac{\lambda^n e^{-\lambda}}{n!} \quad (1)$$

is entirely parametrised by the variable λ , which describes both the mean and variance of the distribution. In actual radio data, some deviations from a perfect Poisson distribution are expected due to clustering and the presence of sources with multiple components. The severity of these effects depends largely on factors such as survey depth, angular resolution, and observing frequency, and is therefore difficult to assess without a thorough analysis of the survey. Such an analysis is beyond the scope of this work, but would follow a structure similar to that of [Siewert et al. \(2020\)](#), who demonstrate for the LOFAR Two-Metre Sky Survey first data release (LoTSS DR1, [Shimwell et al. 2019](#)) that the distribution of source counts converges to a Poisson distribution if applying stricter flux density cuts. For the analysis presented in this work, we assume that the effects of clustering and multi-component sources are negligible on a dipole estimate.

The spectral features and number count relations of typical radio sources make them uniquely suitable for a dipole measurement. For most sources, the dominant emission mechanism at radio wavelengths is synchrotron radiation, the spectral behaviour of which is well described by a power law,

$$S \propto \nu^{-\alpha}, \quad (2)$$

with a characteristic spectral index α . For synchrotron emission, the typical value of α is around 0.75 (e.g. [Condon 1992](#)), and this value has been assumed for most dipole studies at radio wavelengths ([Ellis & Baldwin 1984](#); [Rubart & Schwarz 2013](#); [Siewert et al. 2021](#)). Furthermore, the number density of radio sources follows a power-law relation with respect to the flux density S above which the counts are taken,

$$\frac{dN}{d\Omega}(>S) \propto S^{-x}. \quad (3)$$

The value of x can differ between surveys depending on the choice of flux density cut and frequency of the catalogue, but usually takes values of 0.75–1.0.

In the frame of the moving observer, given a velocity $\beta = v/c$, a systemic Doppler effect shifts the spectra of sources, which affects the flux density of these sources. Additionally, sources are Doppler boosted from the point of view of the moving observer, which further affects the observed flux density of sources. Depending on their angular distance from the direction of motion θ , the flux densities are shifted by

$$S_{\text{obs}} = (1 + \beta \cos \theta)^{1+\alpha} S_{\text{rest}}. \quad (4)$$

Thus, given a flux-limited survey of radio sources, more sources appear above the minimum observable flux in the direction of the motion, and less will appear in the opposite direction. Finally, relativistic aberration caused by the motion of the observer shifts the positions of sources towards the direction of motion, causing a further increase in number counts in the direction of motion,

$$\tan \theta_{\text{obs}} = \frac{\sin \theta_{\text{rest}}}{\beta - \cos \theta_{\text{rest}}}. \quad (5)$$

As the fluxes and positions of the sources are shifted, we observe the dipole as an asymmetry in the number counts of radio sources. Combining these effects to first order in β , and therefore assuming that $v \ll c$, shows the expected dipole amplitude for a given survey:

$$\mathbf{d} = \mathcal{D} \cos \theta, \quad (6)$$

$$\mathcal{D} = [2 + x(1 + \alpha)]\beta. \quad (7)$$

As such, we can directly infer the velocity of the observer by measuring the dipole effect on the number counts of sources. However, this necessitates the assumption that the dipole is entirely caused by the motion of the observer. Given the observed discrepancy between the CMB dipole and the radio dipole, it would be equally appropriate to assume that part of the observed radio dipole is caused by a different (and as-of-yet unknown) effect.

Given a dipole characterised by Eq. (7), different dipole amplitudes \mathcal{D} are expected to be seen depending on the data set being used. Though aberration always has an equal effect, the effect of Doppler shift is determined by the spectral index α of the sources, and both this and the Doppler-boosting effect depend on the flux distribution of sources, which is characterised by the power-law index x . Single values are most often assumed for these quantities, from which the expectation of the dipole amplitude can be derived. These quantities will differ at least between different surveys, and so we derive them for each survey separately. Although the entire flux distribution of any survey cannot be characterised with a single power-law relation, the dipole in number counts is caused by sources near the flux density threshold, making the power-law fit near this threshold the most appropriate choice for deriving x . For the entire range of frequencies considered in this work, we assume a spectral index of $\alpha = 0.75$, considering the synchrotron emission of radio sources which dominates the spectrum of radio sources below 30 GHz (Condon 1992).

3. Dipole estimators

Different types of estimators have been used to measure the cosmic radio dipole, which for the most part yield consistent results. Most commonly used are linear and quadratic estimators (e.g. Singal 2011; Rubart & Schwarz 2013; Siewert et al. 2021). Linear estimators essentially sum up all source positions, and therefore, by design, point towards the largest anisotropy in the data. However, the recovered amplitude from linear estimators is inherently biased. Furthermore, because of the sensitivity of linear estimators to anisotropies, any gaps or systematic effects in the data can introduce biases in the estimate of the dipole direction. To avoid biasing the estimator with respect to dipole direction, a mask must be created such that the map remains point symmetric with respect to the observer, or a ‘masking correction’ must be applied (e.g. Singal 2011; Rubart & Schwarz 2013). Consequently, missing data features such as the Galactic plane must be mirrored to maintain symmetry, removing even more data. The quadratic estimator compares expected number counts with a model, providing a chi-squared test of the data with respect to a model of the dipole (e.g. Siewert et al. 2021). The best-fit dipole parameters are then retrieved by minimising χ^2 . Though the cost is the imposition of a dipole model on the data, the estimate is not biased by the ubiquitous spatial gaps in the data of radio surveys.

Both aforementioned estimators are sensitive to anisotropies in the data introduced by systematic effects. Most commonly,

these systematic effects influence the sensitivity of the survey in different parts of the sky, meaning the most straightforward solution is to cut out all sources below some flux density. However, given this assessment, we might expect information in terms of the sensitivity of the survey in different parts of the sky to help alleviate these biases. The new generation of radio surveys provides catalogues with a wealth of information, including the local root-mean-square (rms) noise, which we exploit in this work.

To improve sensitivity to the dipole, several attempts have been made to combine different radio catalogues. Both Colin et al. (2017) and Darling (2022) worked on combining different radio surveys, using different techniques to deal with systematic differences between the catalogues. Here, we provide an alternative method to combine catalogues for increased sensitivity to the dipole, while accounting for systematic differences between the catalogues.

3.1. Quadratic estimator

To control for differences in pixelation and masking strategies between this and previous works, introducing a new estimator warrants a comparison with known methods of dipole estimation. The quadratic estimator is the closest analogy to the Poisson estimator used here to produce our main results in that it is insensitive to gaps in the data. Its effectiveness and results on multiple large radio surveys are presented in Siewert et al. (2021). The quadratic estimator is based on the Pearson’s chi-squared test, minimising

$$\chi^2 = \sum_i \frac{(n_{i,\text{obs}} - n_{i,\text{model}})^2}{n_{i,\text{model}}}, \quad (8)$$

where the dipole model is written as

$$n_{i,\text{model}} = \mathcal{M}(1 + \mathbf{d} \cdot \hat{\mathbf{n}}_i). \quad (9)$$

Here, the dipole amplitude on a given cell is given by the inner product between the dipole vector \mathbf{d} and the unit vector pointing in the direction of the cell $\hat{\mathbf{n}}_i$, with $\mathbf{d} \cdot \hat{\mathbf{n}}_i = \mathcal{D} \cos \theta_i$. In addition to the dipole vector, the monopole \mathcal{M} is a free parameter, for which the mean value of all cells \bar{n} is a good initial estimate. The χ^2 test is agnostic to the actual distribution of the data, but a dipole model is imposed on the data. This can give rise to misleading results if there are anisotropies in the data – such as those caused by systematic effects – on large enough scales and with sufficiently large amplitude to influence the fit. This can generally be assessed with the reduced χ^2 , which should take a value of around unity if the fit is good.

3.2. Poisson estimator

In the dipole estimators used in previous works, no explicit assumption was made as to the shape of the distribution of sources. A Gaussian distribution can be a valid assumption for a source distribution, although it has an additional degree of freedom compared to Poisson and is less valid if cell counts are low. However, as we do not know a priori how many sources we have, and as we do not commit to a cell size for which we count sources, we choose to assume a Poisson distribution for our cell counts. The Poisson probability density function is given by Eq. (1), and depends on the mean of the distribution λ , which is equal to the monopole \mathcal{M} in the absence of anisotropies. To account for the effect of the dipole, we introduce a dipole model

equivalent to Eq. (9),

$$\lambda(\mathbf{d}, \mathcal{M}) = \mathcal{M}(1 + \mathbf{d} \cdot \hat{\mathbf{n}}). \quad (10)$$

In order to estimate the dipole parameters, we maximise the likelihood, which is given by

$$\mathcal{L}(n|\mathbf{d}, \mathcal{M}) = \prod_i \frac{\lambda(\mathbf{d}, \mathcal{M})^{n_i} e^{-\lambda(\mathbf{d}, \mathcal{M})}}{n_i!}. \quad (11)$$

Maximising the likelihood through posterior sampling has the key advantage of immediately yielding the uncertainties on the derived parameter values. This removes the necessity for null-hypothesis simulations as performed by for example Rubart & Schwarz (2013) and Secrest et al. (2021, 2022). This estimator was recently used by Dam et al. (2022), who showed that it provides tighter constraints on the dipole than previously used methods.

3.3. Poisson-rms estimator

While a survey can be influenced by many different systematic effects, we can assert that the net effect is different sensitivity of the survey at different parts of the sky, leading to anisotropic number counts. Therefore, we assume that all systematic effects that impact source counts in fact influence local noise, thereby causing the source density to vary across the survey. If the survey has sensitivity information in each part of the sky, the impact can be simply modelled by introducing additional variables to the model. So long as a detection threshold is consistently applied to the entire survey, the lower flux density limit will be linearly related to the local rms noise. Consequently, taking into account the dipole and the power law describing number counts, we can model the mean counts in the Poisson estimator as

$$\lambda(\mathbf{d}, \mathcal{M}, \sigma, x) = \mathcal{M} \left(\frac{\sigma}{\sigma_0} \right)^{-x} (1 + \mathbf{d} \cdot \hat{\mathbf{n}}), \quad (12)$$

where σ is the rms noise of the cell, x is the power-law index of the flux distribution, and σ_0 is a reference rms value that scales the power law and explicitly ensures λ is dimensionless. The value of σ_0 does not influence any parameters except for the monopole \mathcal{M} , which will take a value closest to the mean cell count \bar{n} when taking σ_0 as equal to the median rms noise over all cells. As the dipole amplitude depends on the power-law index of the flux distribution near the flux limit, the variation in the rms noise should be small enough that it can be adequately described with a single value. We expect a linear relation between the flux density limit and the local noise, as the detection threshold for most surveys is some multiple of the noise; usually 5σ . Maximising the likelihood given by Eq. (11) while inserting Eq. (12) for λ can therefore yield the best-fit dipole and power-law parameters.

3.4. Multi-Poisson estimator

Hoping to remove any systematic effects stemming from the incomplete sky coverages of individual radio surveys, Darling (2022) combined the Rapid Australian Square Kilometre Array Pathfinder (ASKAP) Continuum Survey (RACS, McConnell et al. 2020) and the VLA Sky Survey (VLASS, Lacy et al. 2020), finding a dipole that, surprisingly, agrees with the CMB in both amplitude and direction, though with large uncertainties. Secrest et al. (2022) note two inherent problems

to this approach. Not only might selecting catalogues at different frequencies select different spectral indices, invalidating the assumption of a common dipole amplitude, but combining the catalogues in such a way ignores systematic effects that can vary between catalogues. Indeed, it is most likely the second factor that plays the most important role, as factors such as observing frequency, array configuration, and calibration can all impact the number counts within a survey in ways that are difficult to predict; this is even more true for independent surveys.

Bearing this in mind, we approach the combination of any two catalogues in a different way. We do not make any attempt to unify the catalogues by matching, smoothing, or creating a common map. Rather, we take both catalogues as independent tracers of the same dipole, allowing the two catalogues to have a different monopole amplitude \mathcal{M} . In the Poisson estimator, we therefore estimate \mathcal{M} separately for each catalogue, turning the likelihood into

$$\mathcal{L}(n_1, n_2|\mathbf{d}, \mathcal{M}_1, \mathcal{M}_2) = \prod_i \frac{\lambda(\mathbf{d}, \mathcal{M}_1)^{n_{1,i}} e^{-\lambda(\mathbf{d}, \mathcal{M}_1)}}{n_{1,i}!} \times \prod_j \frac{\lambda(\mathbf{d}, \mathcal{M}_2)^{n_{2,j}} e^{-\lambda(\mathbf{d}, \mathcal{M}_2)}}{n_{2,j}!}. \quad (13)$$

Any dipole estimates with this likelihood benefit if the (expected) dipole amplitudes of the two catalogues are similar. However, any differences will be absorbed into the overall error budget by virtue of the sampling algorithm. Once again, the likelihood can be maximised through posterior sampling, yielding the best-fit dipole results as well as the monopoles \mathcal{M}_1 and \mathcal{M}_2 for both surveys.

3.5. Priors and injection values

For efficient parameter estimation through posterior sampling, proper priors must be set. While priors on parameters can take on many shapes based on prior knowledge, we take flat priors on all parameters. This only leaves us to define the extent of the probed parameter space, as well as the initial guesses to serve as a starting point for the posterior sampling. In terms of dipole parameters, we separately infer dipole amplitude \mathcal{D} , as well as the right ascension and declination of the dipole direction. We expect the dipole amplitude to take values of around 10^{-2} , but to allow for more variation, the prior on the dipole amplitude is set to $\pi(\mathcal{D}) = \mathcal{U}(0, 1)$. Any point in the sky can represent the dipole direction, and so logically the priors cover the entire sky: $\pi(\text{RA}) = \mathcal{U}(0, 360)$ and $\pi(\text{Dec}) = \mathcal{U}(-90, 90)$. As an initial guess for these parameters, we inject the approximate expected dipole parameters from the CMB: $\mathcal{D} = 4.5 \times 10^{-3}$, $\text{RA} = 168^\circ$, $\text{Dec} = -7^\circ$.

Additionally, the parameters of the distribution of number counts are also estimated. In the basic Poisson case, this is represented by the monopole \mathcal{M} . As the dipole is not expected to meaningfully impact this value, a good initial guess of the monopole is the mean of all cell counts \bar{n} . As the real monopole is likely close to this value, we choose the prior $\pi(\mathcal{M}) = \mathcal{U}(0, 2\bar{n})$. For the Poisson-rms estimator, both a monopole \mathcal{M} and power-law index x are estimated. Before the estimation, we fit a power law to the cell counts in order to obtain initial estimates $\mathcal{M}_{\text{init}}$ and x_{init} , which also function as the initial guesses for these parameters. The initial monopole estimate informs the prior as we use $\pi(\mathcal{M}) = \mathcal{U}(0, 2\mathcal{M}_{\text{init}})$. For the power-law index x , a value of around 0.75–1.0 is always expected, and so we take the prior $\pi(x) = \mathcal{U}(0, 3)$.

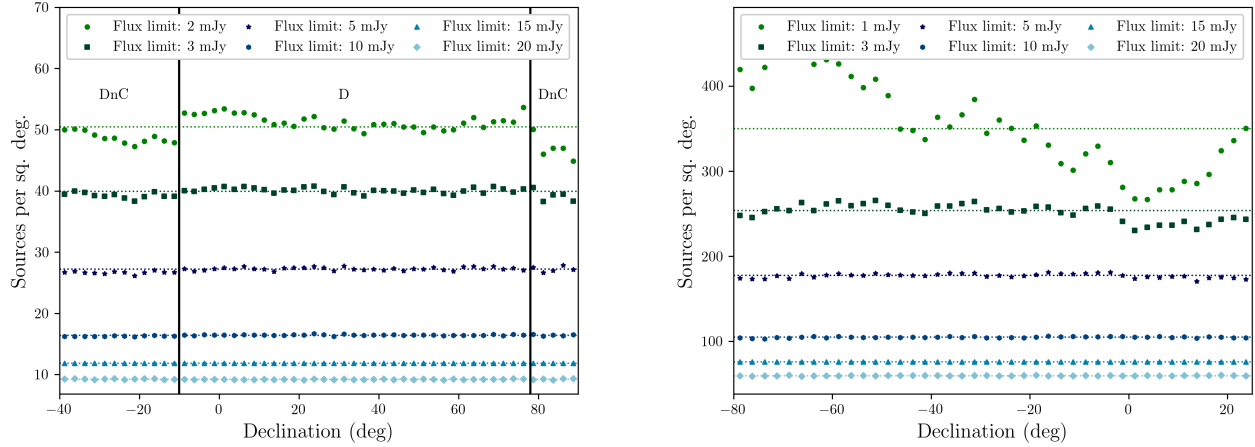


Fig. 1. Source density of NVSS (left) and RACS (right) as a function of declination for different flux density cuts. For the NVSS, we show the boundaries where different array configurations are used. In both cases, the catalogues can be seen to be inhomogeneous at low flux densities.

Table 1. Areas in NVSS masked due to high source density.

Region	RA _{min} (°)	RA _{max} (°)	Dec _{min} (°)	Dec _{max} (°)	Sky area (sq. deg.)
1	82.0	90.0	-7.0	-1.0	47.9
2	49.0	52.0	-39.0	-36.0	7.1
3	185.0	189.0	11.0	14.0	11.7
Total					66.7

4. Data

Given the estimators introduced in Sect. 3, there are a multitude of available radio catalogues to possibly make use of for a dipole estimate. To maintain the focus of our approach, we use NVSS and RACS, two catalogues that cover the full sky when combined, but have been processed in very different ways given the respective eras in which they were produced. Given the introduction of novel estimators, the NVSS is the logical choice for verification, providing a baseline as the most thoroughly studied catalogue in terms of dipole measurements. The choice of RACS for the second catalogue is straightforward; not only does it complement NVSS in terms of sky coverage, but the inclusion of sensitivity information in the catalogue makes it suitable for testing the Poisson-rms estimator described in Sect. 3.3. The complementary sky coverage of the two catalogues also provides the best testing ground for the Multi-Poisson estimator described in Sect. 3.4.

4.1. NVSS

The NVSS (Condon et al. 1998) is one of the most well-studied surveys in terms of dipole measurements, and as such is well suited for verifying novel dipole estimators. It covers the whole sky north of -40° declination, and has a central frequency of 1.4 GHz and an angular resolution of $45''$. The complete catalogue includes the Galactic plane and contains 1 773 484 sources.

An important feature of the NVSS catalogue is that, for observations below a declination of -10° and above a declination 78° , the VLA DnC array configuration was used for observations, while the VLA D configuration was used for the rest of the survey. This affects the number counts at those decli-

nations. The left plot of Fig. 1 shows the source density of NVSS as a function of declination for different flux density cuts. The impact of the different array configurations can be clearly seen at lower flux densities. At only around 15 mJy, the source density becomes homogeneous, and so for an unbiased dipole analysis we choose to exclude all sources with a flux density below 15 mJy, which is a commonly applied flux density cut (e.g. Singal 2011; Siewert et al. 2021). Even after this cut, some areas with significantly high source counts are present in the data. We mask these areas as specified in Table 1 for the dipole estimate.

With a flux density cut at 15 mJy, we fit a power law to the lower end of the flux distribution of sources and find a power-law index of $x = 0.85$. Additionally taking $\alpha = 0.75$ and taking the velocity from the CMB into account (Eq. (7)), this sets the expectation of the dipole amplitude to $\mathcal{D} = 4.30 \times 10^{-3}$.

4.2. RACS

RACS is the first large survey carried out using the Australian Square Kilometre Array Pathfinder (ASKAP), covering the sky south of $+40^\circ$ declination. Observations are carried out with a central frequency of 887.5 MHz and images are smoothed to a common angular resolution of $25''$. The first data release of RACS in Stokes I is described in Hale et al. (2021), and the catalogue used in this work is the RACS catalogue with the Galactic plane removed, containing 2 123 638 sources. Source finding in the images has been done with the Python Blob Detector and Source Finder (PyBDSF, Mohan & Rafferty 2015), which provides a wealth of information on each source. Most importantly, the root-mean-square (rms) noise at the position of each source is present in the noise column of the catalogue, and we use this in the Poisson-rms estimator approach.

The right plot of Fig. 1 shows the source density of RACS as a function of declination for different flux density cuts. Even though there is no change in array configuration as with NVSS, there is a clear gradient in source density, where a greater number of sources are detected at lower declinations. Once again, the catalogue becomes homogeneous at flux densities of around 15 mJy, and so we exclude all sources below this flux density, as with NVSS. Even after this cut, some areas with significantly low source counts are present in the data. These areas appear just above the celestial equator. We mask these areas as specified in Table 2 for the dipole estimate.

Table 2. Areas in RACS masked due to low source density.

Region	RA _{min} (°)	RA _{max} (°)	Dec _{min} (°)	Dec _{max} (°)	Sky area (sq. deg.)
1	357.0	3.0	16.0	22.0	34.0
2	330.0	337.0	16.0	22.0	39.7
3	252.0	261.0	3.0	9.0	53.7
4	184.0	192.0	9.5	15.5	46.8
Total					174.2

With a flux density cut at 15 mJy and the other masks applied, we fit a power law to the lower end of the flux distribution of sources of RACS and find a power-law index of $x = 0.82$. Taking once again $\alpha = 0.75$, this sets the expectation of the dipole amplitude to $\mathcal{D} = 4.24 \times 10^{-3}$.

4.3. Common masks and pixelation

In order to avoid biases to the data, we adopt a masking scheme that uses the same principles for both NVSS and RACS. As mentioned previously, surveys are generally not homogeneous at the lowest flux densities due to variations in the noise, and so we must choose a lower flux density threshold appropriate for the survey. As described above, for both NVSS and RACS we choose a flux density threshold of 15 mJy. To avoid counting Galactic sources or Galactic extended emission, we exclude the Galactic plane. Due to problems with source finding in the Galactic plane, it is already removed in the RACS catalogue, excluding all sources with $|b| < 5^\circ$. For NVSS, overdensities from the Galactic plane extend further out, and we exclude all sources with $|b| < 7^\circ$. Finally, we expect that the brightest sources in a survey push to the limit of the dynamic range, which introduces artefacts around bright sources. Due to differences in source-finding methods between the two catalogues, this increases counts around bright sources in NVSS, but decreases counts around bright sources in RACS. In both cases, the local source density is affected, and therefore we remove all sources within a radius of 0.3° around any source brighter than 2.5 Jy.

After masking data, we use the Hierarchical Equal Area isoLatitude Pixelation (HEALPIX, Górski et al. 2005)² scheme to divide the sky into cells of equal size. HEALPIX allows the flexibility of choosing the number of cells over the whole sky, with the base and minimum value being 12 pixels. The resolution parameter N_{side} determines the number of pixels by $N_{\text{pix}} = 12 \times N_{\text{side}}^2$. We choose two resolutions for our experiment, $N_{\text{side}} = 32$ and $N_{\text{side}} = 64$, which have pixel sizes of $110'$ and $55'$ on a side, respectively. For measuring number counts, each cell holds the number of sources detected within the confines of the cell. To avoid edge effects resulting from pixels covering data only partially, all pixels that have a neighbouring pixel with zero sources are also set to zero. Finally, all pixels with a value of zero are masked to ensure that these pixels are not taken into account during dipole estimation. The HEALPIX maps of NVSS and RACS with $N_{\text{side}} = 32$, with flux density cuts and masks applied, are shown in Fig. 2. For the Poisson-rms estimator, the local noise is determined per HEALPIX cell by taking the median rms of all the sources within the cell.

4.4. Simulations

In addition to the survey data sets of NVSS and RACS, we create a catalogue of simulated sources with a dipole effect to test the validity of the estimators. To do this, we uniformly populate the sky with sources, and assign a rest flux density S_{rest} according to the power law

$$S_{\text{rest}} = S_{\text{low}}(1 - \mathcal{U})^{-1/x}, \quad (14)$$

where \mathcal{U} is a uniform distribution between 0 and 1, S_{low} is the lower flux density limit at which sources are generated, and x the power-law index of the flux density distribution. We transform the rest flux densities and positions of sources by applying relativistic aberration, Doppler shift, and Doppler boost as expressed in Eqs. (4) and (5).

We add Gaussian noise to the flux densities of the sources, generating a larger sample of sources and simulating source extraction by only including sources with $S/N > 5$ in the final catalogue. This naturally adds the effect of Eddington bias to the sample, which is expected to be present in real source catalogues. For a realistic distribution, the local noise variation is taken from the RACS $N_{\text{side}} = 32$ rms map, idealising by assigning the rms of a cell to all sources in that cell. Additionally, we simulate false detections by generating sources with the same flux distribution that are not affected by the dipole. These sources consist of 0.3% of the total catalogue, which is the percentage reported for the RACS catalogues (Hale et al. 2021).

All sources are simulated with a spectral index of 0.75, and the power law used to generate the flux distribution has $x = 1$. We apply the dipole effect assuming the direction derived from the CMB dipole, (RA, Dec) = (170°, -10°), but with an increased velocity of $v = 1107 \text{ km s}^{-1}$ to ensure that a sensitivity to the dipole is reached that is similar to that of NVSS and RACS, with similar monopole values. This sets the expectation of the dipole to $\mathcal{D} = 1.5 \times 10^{-2}$.

5. Results

Using the described estimators, we estimate the dipole parameters for NVSS and RACS. The results are summarised in Table 3, with dipole directions shown in Fig. 3 and dipole amplitudes shown in Fig. 4. To estimate the best-fit parameters, for the quadratic estimator, we minimise χ^2 using LMFIT (Newville et al. 2016). The reduced χ^2 -values are reported in Table 3. For the Poisson estimators, we use the Bayesian inference library BILBY (Ashton et al. 2019), which provides a convenient and user-friendly environment for parameter estimation. We maximise the likelihood using Markov chain Monte Carlo (MCMC) sampling with EMCEE (Foreman-Mackey et al. 2013). The scripts used to obtain these results are available on GitHub³ and an immutable copy is archived in Zenodo (Wagenveld 2023).

To get an indication of how well the distribution fits a Poisson distribution, Pearson's χ^2 can also be used as a Poisson dispersion statistic, defined as

$$\chi^2 = \sum_i \frac{(n_i - \bar{n})^2}{\bar{n}}. \quad (15)$$

As we use the mean of the distribution instead of a model expectation, the resulting χ^2 value indicates the ratio between the

² <http://healpix.sourceforge.net>

³ <https://github.com/JonahDW/Bayesian-dipole>




Fig. 2. Number counts for NVSS (left) and RACS (right) in equatorial coordinates, in $N_{\text{side}} = 32$ HEALPIX maps with masks and a flux density cut of 15 mJy applied.

Table 3. Dipole estimates using the various estimators on the NVSS and RACS catalogues, including a combined estimate using both catalogues.

Survey	Estimator	NSIDE	S_0 (mJy)	N	\mathcal{M} counts/pixel	x	\mathcal{D} ($\times 10^{-2}$)	RA (deg)	Dec (deg)	$\chi^2/\text{d.o.f.}$
Simulation	Poisson	32	1	4 205 065	489.6 ± 0.3	–	31.7 ± 0.1	202 ± 1	-78.5 ± 0.2	37.3
		32	15	434 247	54.5 ± 0.1	–	1.50 ± 0.25	172 ± 10	-17 ± 13	1.01
		32	50	130 150	16.3 ± 0.1	–	2.06 ± 0.60	178 ± 19	-46^{+17}_{-14}	1.02
	Poisson-rms	32	5σ	4 212 472	518.9 ± 0.3	1.000 ± 0.002	1.58 ± 0.08	169 ± 4	-12 ± 4	1.10
	RACS	Quadratic	32	15	451 003	57.4 ± 0.1	–	1.35 ± 0.26	192 ± 12	6 ± 17
64			15	458 152	14.74 ± 0.03	–	1.30 ± 0.24	195 ± 11	4 ± 16	1.11
Poisson		32	15	451 003	56.8 ± 0.1	–	1.41 ± 0.24	193^{+11}_{-10}	4^{+14}_{-15}	1.24
		64	15	458 152	14.19 ± 0.02	–	1.42 ± 0.24	194 ± 10	6 ± 14	1.13
Poisson-rms		32	5σ	2 035 375	253.3 ± 0.2	0.778 ± 0.003	1.62 ± 0.12	210 ± 4	-12^{+7}_{-6}	2.03
		64	5σ	2 068 204	63.8 ± 0.1	0.738 ± 0.003	2.0 ± 0.2	204 ± 5	-40 ± 4	1.49
NVSS	Quadratic	32	15	345 803	41.0 ± 0.1	–	1.35 ± 0.30	145 ± 13	-5 ± 17	1.17
		64	15	352 862	10.66 ± 0.02	–	1.23 ± 0.29	150 ± 13	-10 ± 17	1.09
	Poisson	32	15	345 803	40.4 ± 0.1	–	1.40 ± 0.29	146 ± 12	-5 ± 15	1.26
		64	15	352 862	10.11 ± 0.02	–	1.38 ± 0.29	151 ± 12	-10 ± 15	1.15
NVSS + RACS	Multi-Poisson				\mathcal{M}_1 counts/pixel	\mathcal{M}_2 counts/pixel				
		32	15	796 806	40.4 ± 0.1	56.7 ± 0.1	1.30 ± 0.18	173 ± 9	-1 ± 11	–
		64	15	811 014	10.11 ± 0.02	14.18 ± 0.02	1.29 ± 0.18	175 ± 8	-2 ± 11	–

variance and the mean of the distribution⁴. For a Poisson distribution, the χ^2 value divided by the number of degrees of freedom (d.o.f.), $\chi^2/\text{d.o.f.}$, should therefore be (close to) unity. In the case of the Poisson-rms estimator, number counts are corrected for the derived power law before calculating χ^2 .

5.1. Quadratic and Poisson estimators

Both the quadratic and Poisson estimators are insensitive to gaps in the data, but will be sensitive to inhomogeneous source counts in the data. As such, we perform the flux density cuts on all the data before estimating the dipole. As specified in Sect. 4, we choose a flux density cut of 15 mJy for all catalogues in addition to the other masks described above. As the simulated catalogue uses the local noise information from RACS, the same masking is applied there. This leaves $\sim 3.5 \times 10^5$, $\sim 4.5 \times 10^5$, and $\sim 2.2 \times 10^6$ sources for NVSS, RACS, and the simulated catalogue, respectively. The resulting number counts for NVSS and RACS are shown in Fig. 2. Along with estimating the dipole amplitude and direction, we estimate the monopole \mathcal{M} . For NVSS and RACS, these parameters

are estimated using HEALPIX maps of both $N_{\text{side}} 32$ and 64; the results of the different cell sizes are shown in Table 3.

The best-fit parameters for the simulated data set are shown in Table 3, that is, for a low threshold of 1 mJy, the common threshold of 15 mJy, and a high threshold of 50 mJy. As the noise variation of the simulated catalogue is based on RACS, the 15 mJy threshold should be appropriate for obtaining a good estimate of the injected dipole. The low threshold shows that the dominant anisotropy from the RACS noise, which dominates the dipole by three orders of magnitude, is mostly a declination effect, but a smaller effect in right ascension is also observed. With the 15 mJy threshold, the injected values for the dipole are retrieved within the uncertainties. To see how the estimator reacts to a lack of sources, for the 50 mJy threshold, the required number counts are not reached for a 3σ measurement of the dipole amplitude. However, this does not introduce a bias, as values still match the injected values, albeit with large uncertainties.

Comparing results between the quadratic and basic Poisson estimators on NVSS and RACS, values match within the uncertainties for all estimated parameters. For NVSS, the results of the quadratic estimator and Poisson estimator match those of Siewert et al. (2021) for a flux density cut of 15 mJy in terms of dipole amplitude, but the direction is slightly offset ($\Delta\theta \sim 20^\circ$). This is caused by a difference in masking strategy; as shown

⁴ Peebles (1980) uses this measure as a clustering statistic of the large-scale structure, and specifically to define the number of objects per cluster.




Fig. 3. Best-fit dipole directions for the Poisson estimator of NVSS (blue), RACS (red), RACS with rms power law (green), and NVSS+RACS (purple) compared with the CMB dipole direction (black star). Different transparency levels represent 1σ , 2σ , and 3σ uncertainties. In all cases, results from the $N_{\text{side}} = 32$ HEALPix map are shown.




Fig. 4. Best-fit dipole amplitudes with 3σ uncertainties for the Poisson estimator of NVSS (blue), RACS (red), RACS with rms power law (green), and NVSS+RACS (purple), compared with an expected CMB dipole amplitude of $\mathcal{D} = 4.5 \times 10^{-3}$ (black line). Both results from the $N_{\text{side}} = 32$ and $N_{\text{side}} = 64$ HEALPix maps are shown.

by Siewert et al. (2021), different masks yield different dipole parameters. The low χ^2 values for the Poisson estimators indicate that a Poisson assumption is in line with the expected distribution of source counts.

As the results in Table 3 indicate for both quadratic and Poisson estimators, the results between $N_{\text{side}} = 32$ and $N_{\text{side}} = 64$ pixel sizes agree with each other within the uncertainties. Furthermore, the dipole amplitudes of NVSS and RACS also agree with each other within the uncertainties, with the dipole amplitude from RACS being slightly higher. The dipole directions between RACS and NVSS are somewhat misaligned ($\Delta\theta \sim 50^\circ$), though both align with the CMB dipole direction within 3σ , that is, at $\Delta\theta \sim 20^\circ$ and $\Delta\theta \sim 30^\circ$ for NVSS and RACS, respectively (see also Fig. 3). Figure 4 shows the amplitudes of the results of the Poisson estimator on NVSS and RACS, including uncertainties. In all cases, the amplitude of the dipole is 3–3.5 times higher than the dipole amplitude expectation from the CMB. For NVSS, the result is at 3.4σ significance and for RACS at a significance of 4.1σ .

5.2. Poisson-rms estimator

As described in Sect. 3.3, we aim to account for the variation in source counts across the survey by assuming these are described by the rms noise of the images. For this estimator, we do not apply the flux density cut, and instead fit a power law that relates the rms of a cell to the number counts in that cell. The rms of the survey is not available for NVSS, but is present in the RACS catalogue for each source individually. We obtain the rms of a cell by taking the median rms value of all sources within it. We take the median rms of all cells as the reference rms, which is $\sigma_0 = 0.33 \text{ mJy beam}^{-1}$. The HEALPix maps of source counts and median rms per cell for RACS are shown in Fig. 5, showing the variation of rms and source counts across the survey. Along with estimating the dipole parameters, the monopole \mathcal{M} and power-law index x are estimated as well. For RACS, these parameters are estimated using HEALPix maps of both $N_{\text{side}} = 32$ and $N_{\text{side}} = 64$.

The parameters for the simulated data set are estimated and shown in Table 3. The noise variation of the simulated catalogue is based on the $N_{\text{side}} = 32$ RACS rms map shown in Fig. 5, which in this case means that the rms map is a perfect representation of the noise in the catalogue. The rms estimator retrieves the injected dipole parameters, with a much higher significance than the standard Poisson estimator.

The results for RACS are shown in Table 3, showing a rather large discrepancy between the two pixel scales, and with respect to other results as well. In both cases, the dipole amplitude is increased with respect to the quadratic and basic Poisson estimators, and the direction is no longer agreeing with the direction of the CMB dipole. The $N_{\text{side}} = 32$ map seems to be less affected than the $N_{\text{side}} = 64$ map, but in both cases the dipole direction is further away from the CMB dipole direction, with $\Delta\theta \sim 40^\circ$ and $\Delta\theta \sim 45^\circ$ separation for the $N_{\text{side}} = 32$ and $N_{\text{side}} = 64$ maps, respectively. Especially striking is the recovered dipole direction of the $N_{\text{side}} = 64$ map, which is at a declination of -40 deg. This retrieved dipole direction aligns towards the anisotropy retrieved in the simulated data with the 1 mJy flux density threshold. Rather than pointing to an additional systematic effect that is not modelled by the local rms, it is therefore more likely that the median rms noise per cell does not adequately represent the noise variation observed in the catalogue.

To further investigate these results, the power-law fits to the cells are shown in Fig. 6, indicating that both power laws are a good fit to the distribution. For the $N_{\text{side}} = 64$ map, the relation fits less well to the cells with lower number counts, possibly indicating that the power-law assumption breaks down for these cells. One effect that can contribute to this is that, at such low number counts, the median rms will be a less robust measure of the local noise. As is the case for the other RACS results, there is a misalignment in right ascension that is even more pronounced here (see also Fig. 3). As seen in Fig. 4, the dipole amplitude is also increased. For the $N_{\text{side}} = 32$ map, the dipole amplitude is 3.8 times higher than the CMB expectation with a formal significance of 10σ , and for the $N_{\text{side}} = 64$ map the dipole amplitude is 4.7 times higher with a formal significance of 8σ .

5.3. Combining RACS and NVSS

Following the procedure laid out in Sect. 3.4, we obtain a combined estimate of the dipole parameters of NVSS and RACS, assuming a common dipole amplitude but independent monopole amplitudes. Although we show that slightly different dipole amplitudes are to be expected between the catalogues,




Fig. 5. Number counts (left) and median rms (right) for RACS in equatorial coordinates, in $N_{\text{side}} = 32$ HEALPIX maps, with no flux density cuts applied.

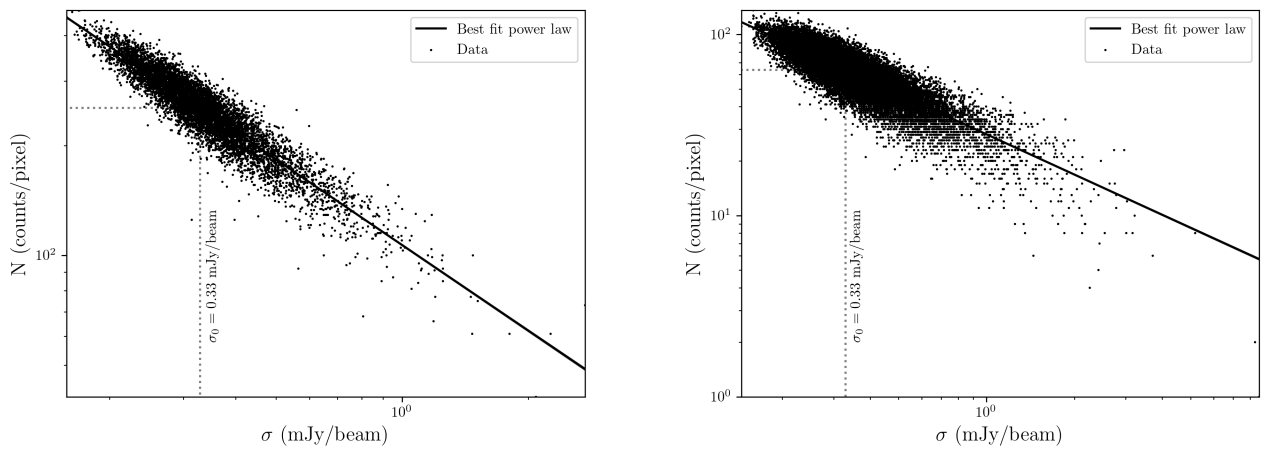


Fig. 6. Cell counts of the $N_{\text{side}} = 32$ (left) and $N_{\text{side}} = 64$ (right) HEALPIX maps with no flux density cuts applied, as a function of the median rms of the pixels, along with the best-fit power-law model (black solid line). The determined σ_0 is indicated by the dashed vertical line; it intercepts the power law at the best-fit monopole value indicated by the horizontal dashed line.

the degree of this difference depends entirely on the degree to which the inferred dipole is kinetic. Nevertheless, because of the nature of parameter estimation with MCMC, any differences in dipole amplitude between the catalogues will be absorbed into the overall uncertainty of the estimated parameters. In terms of monopole, there is no question as to the difference between the catalogues, as from the estimates of the individual catalogues – even with the same cut in flux density – there are large differences in source density.

Table 3 shows the results of the dipole parameters for the combined estimate of NVSS + RACS. Whereas the dipole directions for the individual catalogues were misaligned with the CMB dipole direction, the combined estimate favours a dipole direction that is perfectly aligned ($\Delta\theta = 4^\circ$, see Fig. 3) with that of the CMB dipole. In line with this finding, the dipole amplitude is reduced with respect to either of the individual catalogues; however, is still in tension with the CMB dipole. The dipole amplitude is three times higher than the CMB expectation, with a significance of 4.8σ for both the $N_{\text{side}} = 32$ and $N_{\text{side}} = 64$ maps. If we base our belief in a dipole result on its agreement with the CMB dipole in terms of direction, then this is the most significant and reliable result we obtain in this work. It is furthermore the most significant result obtained with radio sources to date, matching the significance of the dipole estimate with WISE AGN from Secrest et al. (2022), although less significant than the joint WISE+NVSS result from the same work.

6. Discussion

The results presented here demonstrate the potential of our introduced estimators and present at the very least an alternative method of making dipole measurements in present and future surveys. Figure 7 shows the results from this work compared to the most recent results from various surveys taken from Siewert et al. (2021) and Secrest et al. (2022) in terms of dipole amplitude. The NVSS is consistent across all works, as it is with our results. The results obtained here agree within uncertainties with those from Secrest et al. (2022), with the exception of RACS-rms and the WISE measurements. The same goes for the findings of Siewert et al. (2021), with the exception of the TGSS result. Though the extremely high amplitude of TGSS is attributed to a frequency dependence of the dipole in Siewert et al. (2021), the WISE result from Secrest et al. (2022) does not follow the fitted trend. Secrest et al. (2022) suggest that the TGSS result might deviate due to issues in flux calibration. As we have measurements at different frequencies, we can tentatively check whether results match up with the frequency evolution model of the dipole amplitude from Siewert et al. (2021), which predicts $\mathcal{D} = 2.3 \times 10^{-2}$ at the RACS frequency of 887 MHz. Our RACS result for a flux cut of 15 mJy, which agrees with NVSS, does not follow the trend predicted by the model; however, the 150 mJy TGSS flux density cut made in Siewert et al. (2021) corresponds to a 40 mJy flux cut in RACS (assuming $\alpha = 0.75$). Applying this flux density cut using the




Fig. 7. Dipole amplitudes with 3σ uncertainties compared to the amplitude expected from the CMB from this work and to results from Siewert et al. (2021) and Secrest et al. (2022). The results from the different works are separated by horizontal dashed lines, showing the results from this work at the bottom.

$N_{\text{side}} = 32$ RACS map, the inferred dipole direction shifts by $\Delta\theta = 8.8^\circ$, and the dipole amplitude increases to $\mathcal{D} = (1.94 \pm 0.37) \times 10^{-2}$. As such, our results cannot rule out the frequency dependence predicted by Siewert et al. (2021), but the obtained results from WISE AGN (Secrest et al. 2021, 2022; Dam et al. 2022) provide a strong argument against it. Though our results are consistent with the literature, as with many works concerning the dipole, their validity and that of the methods require further examination.

6.1. The Poisson solution

Though we show results here that are both internally consistent and consistent with other dipole estimates, the choice of a Poisson estimator might seem like an unnecessary constraint on the data; after all, the quadratic estimator shows an adequate performance and does not suffer any loss in precision compared to the Poisson estimator. Table 3 lists the χ^2 values for the obtained results, defined by Eqs. (8) and (15) for the quadratic and Poisson estimators, respectively. As the quadratic estimator is minimised for a Gaussian distribution with mean equal to the variance, the quadratic and basic Poisson estimators are expected to provide similar values. This is indeed the case for the results in Table 3, both in the estimated parameters and χ^2 values.

The value of the Poisson assumption becomes readily apparent when extending the parameter space, as we do when taking into account the rms power-law relation. The main feature of a Poisson distribution is that one parameter is necessary to describe it, λ , which is both the mean and variance of the distribution. This is a strict requirement on a distribution, allowing more freedom in other parameters which would otherwise be degenerate with the parameters of the distribution. This means that fitting the rms power law does not work with a quadratic estimator for example; indeed, this latter, though minimised by a distribution with mean equal to the variance, still allows for a wider Gaussian distribution. As seen in Fig. 8, the distribution of number counts without any flux density cut applied resembles a Gaussian distribution, which is much wider than a Poisson distribution with the same mean. However, the quadratic estimator does allow such a wide distribution, and therefore will not converge on a solution that transforms this distribution to a Pois-

son distribution. Herein lies the power of the Poisson estimator, which makes modelling and fitting of systematic effects in the data a viable alternative to cutting and masking data. Nevertheless, one drawback is that it is imposing a Poisson distribution on the data, which can lead to spurious results if improperly applied.

Table 3 lists the $\chi^2/\text{d.o.f.}$ values of the Poisson rms estimator after correction for the derived power law. The difference in distributions can be appreciated in Fig. 8, which shows the distributions of the cell counts of RACS without any flux density cut applied, along with the same distribution corrected for the rms power law that has been fit to the data. The uncorrected counts have a much wider distribution, which is clearly not Poisson, with $\chi^2/\text{d.o.f.} = 13.28$ for the $N_{\text{side}} = 32$ map and $\chi^2/\text{d.o.f.} = 4.49$ for the $N_{\text{side}} = 64$ map. The corrected counts resemble a Poisson distribution more closely, with $\chi^2/\text{d.o.f.} = 2.03$ for the $N_{\text{side}} = 32$ map and $\chi^2/\text{d.o.f.} = 1.49$ for the $N_{\text{side}} = 64$ map, but $\chi^2/\text{d.o.f.}$ values indicate variance is still too large for a Poisson distribution, signifying that some residual effect has not been modelled by the estimator.

As such, the performance of the Poisson-rms estimator still leaves some questions to be answered. The assumption that source counts are related to sensitivity via a power law might carry a flaw, though there can be a number of possible reasons for this: (i) the median rms is not the best representation of the sensitivity of the survey in a given cell; (ii) the sensitivity only properly represents source counts down to some limit; and (iii) not all systematic effects equally impact source counts as well as sensitivity. These factors require further examination in the future, but remarkable already are the results when compared to the other RACS results. It is clear that this Poisson-rms estimator shows promise even in its basic form, and can be used as an additional test of the data for any survey that has information on the local rms. Furthermore, due to its flexibility, additional effects once characterised can easily be modelled and taken into account by the estimator.

6.2. Residual anisotropies in the data

In dipole measurements and other statistical studies that require large amounts of data to retrieve a statistically significant measure, it can be difficult to visually assess whether any one fit adequately describes the data. After all, we impose a model on the data to which the fit is restricted. For a rudimentary visual verification of whether or not the data follow the expected relations, we employ the hemisphere method used by Singal (2021). This method assumes that the direction of the dipole is already known, leaving the dipole amplitude as a function of angular distance from the dipole direction, $\mathcal{D}_\theta = \mathcal{D} \cos \theta$, as the only free parameter. To reach statistically significant number counts, the sky is divided into two hemispheres: hemisphere N_1 with all sources between θ and $\theta + \pi/2$, and hemisphere N_2 with all sources between $\theta + \pi/2$ and $\theta + \pi$. The dipole amplitude as a function of θ is then written as

$$\mathcal{D}_\theta = \frac{N_1(\theta) - N_2(\theta)}{\frac{1}{2}[N_1(\theta) + N_2(\theta)]}. \quad (16)$$

We determine and plot the hemisphere results for NVSS and RACS assuming the results obtained from the Poisson estimators for the individual catalogues. The hemisphere relation for NVSS is presented in Fig. 9, which shows the data following the expected dipole curve except for the hemispheres closest to the dipole direction; these data reveal an increased anisotropy. This is more pronounced in the $N_{\text{side}} = 64$ map, where both $\theta = 0^\circ$ and

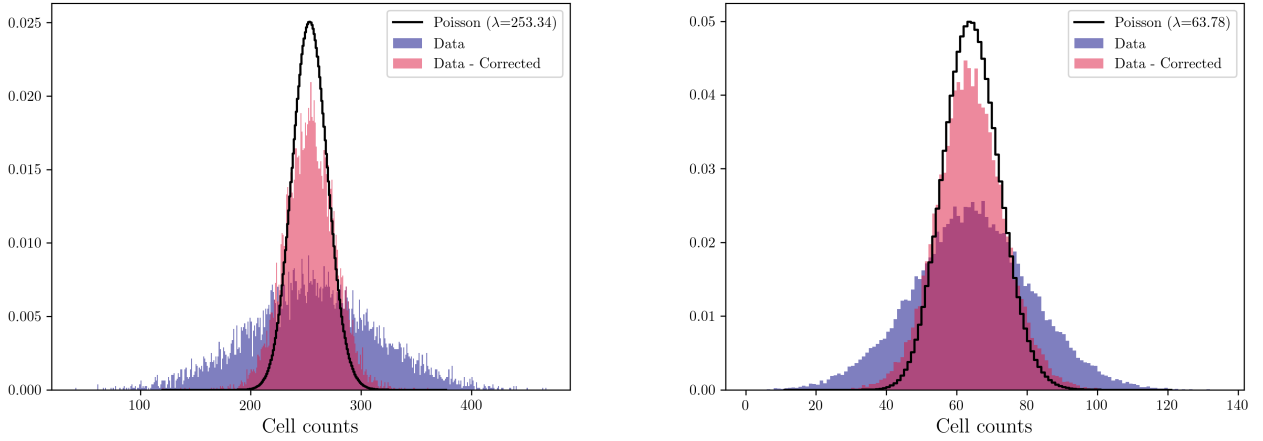


Fig. 8. Cell count distribution for $N_{\text{side}} = 32$ (left) and $N_{\text{side}} = 64$ (right) RACS maps without any flux density cuts applied. Raw counts are shown (blue histogram) alongside the counts corrected for the power-law fit (red histogram). A Poisson distribution with a λ equivalent to the estimated monopole amplitude is also shown (black histogram).

$\theta = 10^\circ$ hemispheres show significantly increased counts compared to the expectation from the dipole model. In the $N_{\text{side}} = 32$ map, only the $\theta = 0^\circ$ hemisphere shows increased counts, with all other points following the dipole model within uncertainties. This points to a residual anisotropy left in the data that has not influenced the overall fit.

For RACS, the hemisphere relations for both the basic Poisson estimator and the rms Poisson estimator are shown in Fig. 10. There is a residual anisotropy in RACS at 40° – 60° from the dipole direction for both $N_{\text{side}} = 32$ and $N_{\text{side}} = 64$ maps that stands out immediately, especially in the results for the basic Poisson estimator. In the case of the basic Poisson estimator, as with NVSS, the fit is unaffected. However, this anisotropy might have had a significant impact on the rms Poisson estimator for the $N_{\text{side}} = 64$ map, as the dipole direction estimated from that map is 47° offset from the direction of the basic Poisson estimator, coinciding perfectly with the anisotropy seen at that angle. Indeed, the data for both $N_{\text{side}} = 32$ and $N_{\text{side}} = 64$ maps for the rms Poisson estimate agree well with the exception of the points closest to the dipole direction, which in the case of the $N_{\text{side}} = 64$ map are dominating the fit.

Finally, we investigate the possibility that residual systematic effects are present due to Galactic synchrotron. Secrest et al. (2022) use the de-stripped and source-subtracted Haslam et al. (1982) 408 MHz all-sky map from Remazeilles et al. (2015) to mask pixels bright in Galactic synchrotron. To investigate if our results are impacted by Galactic synchrotron, we cross-correlate our number count maps with the Remazeilles et al. (2015) map. In all cases, no significant correlation is found ($|\rho| < 0.05$), showing that Galactic synchrotron is not present as a residual systematic effect in the data.

6.3. Combining and splitting catalogues

As shown in Sect. 5.3, combining catalogues as independent tracers of the dipole can provide a more robust measurement of the dipole, and in the case of NVSS and RACS, result in a dipole that matches the direction of the CMB dipole remarkably well, with a dipole amplitude of 2.5 times the CMB expectation. The justification for this approach is that radio data are a complex product and are sufficiently difficult to homogenise over a full survey internally, not to mention between surveys. Field of view, frequency, array configuration, calibration, imag-




Fig. 9. Dipole amplitude of NVSS as a function of angular distance to the dipole direction, assuming a best-fit dipole direction from the basic Poisson estimator, for both $N_{\text{side}} = 32$ (blue) and $N_{\text{side}} = 64$ (red) maps. Alongside the data, the corresponding models are plotted (solid lines), as well as the expected model from the CMB dipole (black dashed line).

ing and source finding are all factors to consider when assessing the source counts in a given survey. For a key example of how these factors can influence source counts, we need to look no further than NVSS, which has been observed with two different array configurations at certain declination ranges. This is a systematic effect that produces different source densities depending on the configuration, something that can be plainly seen in the left plot of Fig. 1.

To demonstrate the potential of the Multi-Poisson estimator beyond combining independent catalogues, we repeated the dipole estimates for NVSS, lowering the minimum flux density in several steps. We then split the NVSS D and DnC configurations into separate catalogues, and repeated the experiment. The results are shown in Table 4 and Fig. 11, showing the effect of splitting the configuration on the dipole estimates. Immediately, we can see that, though the dipole estimate in both cases remains consistent in terms of right ascension, the separation of the catalogue produces wildly different results for flux density cuts below 15 mJy, both in terms of dipole amplitude and declination of the dipole direction. This result is expected to some degree, as the difference in source density between the D and DnC

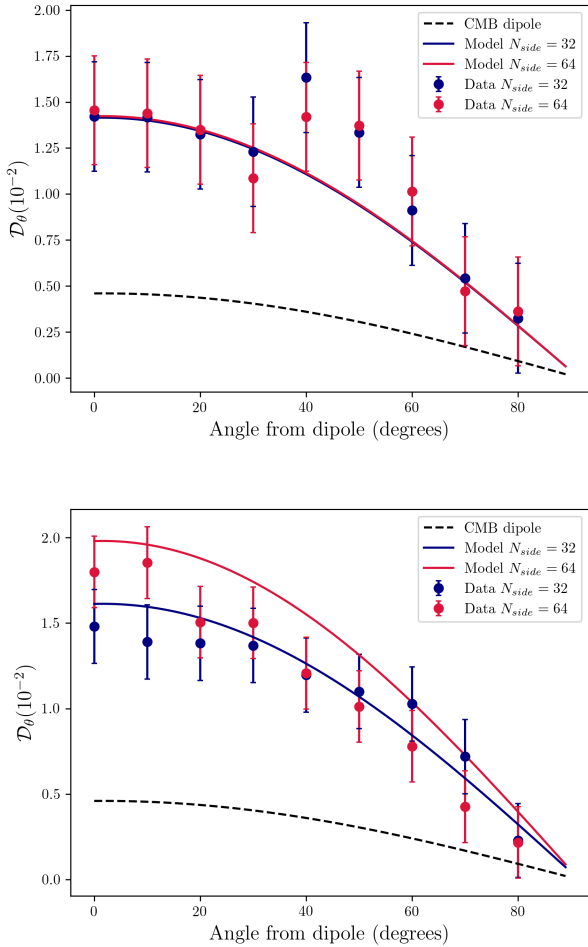


Fig. 10. Dipole amplitude of RACS as a function of angular distance to the dipole direction, assuming a best-fit dipole direction from the basic Poisson estimator (top) and Poisson rms estimator (bottom). Results for both $N_{\text{side}} = 32$ (blue) and $N_{\text{side}} = 64$ (red) maps are shown. Alongside the data, the corresponding models are plotted (solid lines), as well as the expected model from the CMB dipole (black dashed line).

configurations is expected to produce an anisotropy in the north–south direction, which is largely alleviated (although not entirely) with the split in configurations. This is not only reflected in the declination of the dipole direction, but in the dipole amplitude as well, which is seen to decrease with more samples in the case of split configurations. Though the anisotropy in the declination is alleviated by this, another anisotropy in right ascension seems to start dominating at lower flux densities, dragging the dipole direction 43° from the CMB dipole direction.

Though it seems that the different NVSS configurations produce an anisotropy that is on a similar level to other anisotropies related to incompleteness of the catalogue, and cannot therefore be reliably used to completely account for the systematic effects in the survey that appear when employing lower flux density thresholds, these results show that were such systematic effects to dominate the catalogue, restructuring the problem to consider these as multiple independent catalogues can produce sensible results. Our analysis is furthermore a useful test of the chosen flux density threshold, as for an appropriately chosen flux density threshold the results between the full catalogue and split catalogue should be consistent. The results here show that a flux density threshold of 15 mJy is indeed appropriate for the NVSS.

Such an approach might become particularly relevant given that VLASS also uses different array configurations, which can be taken into account when estimating a dipole in the same manner as we have done for NVSS. Even for RACS, a prominent feature is a dependence of source counts on declination; though the exact mechanism is unclear, one possibility is related to the point spread function, as the UV-coverage of the array evolves with declination.

6.4. Combining catalogues and cosmological considerations


The combined dipole estimate of RACS and NVSS makes a compelling case for combining more probes of the cosmic dipole to increase sensitivity. The approach used here ostensibly carries less caveats than previous works, foregoing source matching, frequency scaling, subsampling, and weighting schemes, which can all introduce additional uncertainties. This carries with it a reduction in formal uncertainty, though there remain some factors regarding the nature of the dipole that can limit the approach. As the nature of the excess amplitude of the radio dipole is currently unknown, the approach of combining catalogues, even if done perfectly, carries an additional uncertainty. In Sect. 4, the expected kinematic dipole amplitudes for both NVSS and RACS are computed and are found to be nearly identical. This in large part justifies the obtention of a combined estimate of the catalogues; however, if we were to combine catalogues where the expected dipole amplitudes differed (e.g. when combining catalogues from multiple wavelengths), the approach would not be able to produce a reliable result without knowing the nature of the radio dipole. As it stands, we have a kinematic expectation of the radio dipole derived from the CMB. Given the velocity of the observer, this dipole is determined by the spectral index and flux distribution of the ensemble of sources. However, because of the measurement method used here, the excess dipole has an unknown origin, be it either kinetic or due to some entirely different effect.

Given the results obtained so far in this work and the literature, one could also assume an effect that somehow boosts the observed dipole with respect to the CMB dipole, or an additional anisotropy that is simply added to the kinematic dipole. Therefore, should we wish to combine for example the sample of WISE AGN, which has an expected dipole amplitude of $\mathcal{D} = 7.3 \times 10^{-3}$ (Secrest et al. 2022), with for example NVSS, a combined estimate would have to assume one of these models. Secrest et al. (2022) find that after removing the CMB dipole, assuming it is purely kinematic, the residual dipoles between NVSS and WISE agree with each other, favouring the interpretation of an intrinsic dipole anisotropy in the CMB rest frame. The results we obtain for NVSS and RACS also support this interpretation, the residual dipole amplitudes being $\mathcal{D} = (0.97 \pm 0.30) \times 10^{-2}$ and $\mathcal{D} = (0.99 \pm 0.24) \times 10^{-2}$, respectively. However, the expected dipoles of the catalogues are too similar to rule out other interpretations.

Furthermore, depending on which of the models presented above is the most accurate, survey design can have a profound impact on the measured cosmic radio dipole. The largest impact will be in the detected source populations and their redshift distributions. Naturally, going to optical or infrared wavelengths will yield different source populations that possibly trace the dipole differently, but even amongst radio surveys the detected source population will depend on the survey details. Radio surveys must be designed with a balance of depth and sky coverage, and so we imagine a scenario where the number of sources detected will stay constant over the survey due to this balance,

Table 4. Dipole estimates for NVSS using different flux density cuts, separating D and DnC configurations.

Catalogue	S_0 (mJy)	N	\mathcal{M}_1 counts/pixel	\mathcal{M}_2	\mathcal{D} ($\times 10^{-2}$)	RA (deg)	Dec (deg)
NVSS	15	345 803	40.4 ± 0.1	–	1.40 ± 0.29	146 ± 12	-5 ± 15
	10	480 446	56.1 ± 0.1	–	1.38 ± 0.26	133 ± 11	31 ± 12
	5	795 135	92.6 ± 0.1	–	1.92 ± 0.23	131 ± 9	53 ± 6
NVSS D + NVSS DnC	15	333 046	40.5 ± 0.1	40.4 ± 0.1	1.44 ± 0.31	148 ± 12	-9 ± 21
	10	463 398	56.4 ± 0.1	55.7 ± 0.2	1.23 ± 0.27	136 ± 12	-8_{-20}^{+22}
	5	767 832	93.5 ± 0.2	91.3 ± 0.2	1.18 ± 0.20	133 ± 9	-3_{-16}^{+18}

**Fig. 11.** Best-fit dipole directions with 1σ uncertainties for the complete NVSS (red) and split NVSS D + DnC (blue) catalogues, for lower flux density thresholds of 5, 10, and 15 mJy. The CMB dipole direction is indicated with a black star.

leaving the significance of a dipole measurement unchanged. A shallow but large-sky-coverage radio survey will mostly detect AGN with a peaked redshift distribution, whereas a deep radio survey with limited sky coverage will probe most of the AGN population over all redshifts, as well as star-forming galaxies, which have a different redshift distribution from AGN altogether. Most surveys that have been used for dipole estimates fall into the first category, as proper coverage along the dipole axis is necessary. However, a survey falling into the second category would have the potential to differentiate between the possible models of the radio dipole we have laid out. The MeerKAT Absorption Line Survey (Gupta et al. 2016), consisting of sparsely spaced deep pointings homogeneously distributed across the sky, provides a good candidate for a survey falling into this second category, and therefore might provide more insight into the processes driving the anomalous amplitude of the radio dipole.

7. Conclusion

In this work, we present a set of novel Bayesian estimators for the purpose of measuring the cosmic radio dipole with the NVSS and RACS catalogues. Based on the assumption that counts-in-cell of radio sources follow a Poisson distribution, we construct estimators for the cosmic radio dipole based on Poisson statistics. To provide a means of for comparison, we include a quadratic estimator, which has been used in a number of previous dipole studies. We furthermore construct two extensions of

the basic Poisson estimator to attempt to account for systematic effects in the respective catalogues. Firstly, we consider that if sensitivity information is present in the catalogue, this can be directly related to the local number density, assuming that systematic effects merely modify the local sensitivity of the catalogue. The local sensitivity and number counts are assumed to be related by a power law, the parameters of which can be estimated. We extend the Poisson estimator to address this. Secondly, we construct an extension to the Poisson estimator that can be given multiple separate catalogues –assuming that the catalogues trace the same dipole– and produces a combined estimate.

We obtain best-fit parameters for the cosmic radio dipole using χ^2 minimisation for the quadratic estimator and using maximum likelihood estimation for the Poisson estimators. To discretise the sky, we use HEALPIX, producing maps with both $N_{\text{side}} = 32$ and $N_{\text{side}} = 64$. We verify that the quadratic estimator and basic Poisson estimator yield similar results, and that furthermore results between the pixel scales are consistent. We use the Poisson rms estimator on RACS while not using any cut in flux density to estimate the dipole parameters along with the parameters for the rms power law. The increased number counts greatly increase the precision of the estimate, but the results somewhat diverge from the dipole estimates produced by the basic estimators. Whether this difference is a genuine product of the data or a flaw in (assumptions of) the estimator is not perfectly understood, but the initial results are still promising given that the entire catalogue of sources is used. We finally use a Poisson estimator for multiple catalogues on NVSS and RACS and obtain a dipole estimate that perfectly aligns with the CMB dipole in terms of direction, but has an amplitude that is three times as large with a significance of 4.8σ . Given the dipole estimates obtained from the individual catalogues, this result is in line with expectations for a combination of the two catalogues, and can therefore be seen as the most reliable and significant result obtained here.

We explore the possibility of splitting up a catalogue and using the Poisson estimator for multiple catalogues to estimate the dipole as if on two independent catalogues. We use this method on the NVSS, which has been observed with two different array configurations, introducing an artificial north–south anisotropy in the catalogue. We treat these array configurations as separate catalogues and repeat the dipole estimate, and go down to lower flux density limits than with the basic Poisson estimator. We see that while using the whole NVSS, a north–south anisotropy starts to dominate the estimate at lower flux densities; separating the configuration largely mitigates this effect. As a result, this allows us to lower the flux density cut, increasing number counts and thus increasing the significance of

the dipole estimates. This approach may work well on catalogues such as VLASS, which also uses different array configurations in different parts of the sky. The presented estimator may provide the potential to combine a larger variety of catalogues, but the extent to which this can be done depends in large part on the nature of the excess dipole. With an increasing array of probes at various wavelengths, sky coverage, and depth, reaching the necessary sensitivity to detect the dipole, the nature of this dipole may well soon be discovered, given the different populations of sources that can be probed.

Acknowledgements. We thank the anonymous referee for their useful comments and feedback on the text. J.D.W. acknowledges the support from the International Max Planck Research School (IMPRS) for Astronomy and Astrophysics at the Universities of Bonn and Cologne. Some of the results in this paper have been derived using the healpy (Zonca et al. 2019) and HEALPIX (Górski et al. 2005) packages. This research has made use of TOPCAT (Taylor 2005), LMFIT (Newville et al. 2016), BILBY (Ashton et al. 2019), and EMCEE (Foreman-Mackey et al. 2013).

References

- Aghanim, N., Akrami, Y., Arroja, F., et al. 2020, *A&A*, **641**, A1
- Ashton, G., Hübner, M., Lasky, P. D., et al. 2019, *ApJS*, **241**, 27
- Blake, C., & Wall, J. 2002, *Nature*, **416**, 150
- Böhme, L., Schwarz, D. J., de Gasperin, F., Röttgering, H. J. A., & Williams, W. L. 2023, *A&A*, **674**, A189
- Colin, J., Mohayaee, R., Rameez, M., & Sarkar, S. 2017, *MNRAS*, **471**, 1045
- Condon, J. J. 1992, *ARA&A*, **30**, 575
- Condon, J. J., & Ransom, S. M. 2016, *Essential Radio Astronomy* (Princeton: Princeton University Press)
- Condon, J. J., Cotton, W. D., Greisen, E. W., et al. 1998, *AJ*, **115**, 1693
- Dalang, C., & Bonvin, C. 2022, *MNRAS*, **512**, 3895
- Dam, L., Lewis, G. F., & Brewer, B. J. 2022, ArXiv e-prints [arXiv:2212.07733]
- Darling, J. 2022, *ApJ*, **931**, L14
- Ellis, G. F. R., & Baldwin, J. E. 1984, *MNRAS*, **206**, 377
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
- Guandalin, C., Piat, J., Clarkson, C., & Maartens, R. 2022, ArXiv e-prints [arXiv:2212.04925]
- Gupta, N., Srikanand, R., Baan, W., et al. 2016, *Proceedings of MeerKAT Science: On the Pathway to the SKA, 25–27 May, 2016 Stellenbosch, South Africa (MeerKAT2016)*, 14
- Hale, C. L., McConnell, D., Thomson, A. J. M., et al. 2021, *PASA*, **38**, e058
- Haslam, C. G. T., Salter, C. J., Stoffel, H., & Wilson, W. E. 1982, *A&AS*, **47**, 1
- Intema, H. T., Jagannathan, P., Mooley, K. P., & Frail, D. A. 2017, *A&A*, **598**, A78
- Lacy, M., Baum, S. A., Chandler, C. J., et al. 2020, *PASP*, **132**, 035001
- Mauch, T., Murphy, T., Buttery, H. J., et al. 2003, *MNRAS*, **342**, 1117
- McConnell, D., Hale, C. L., Lenc, E., et al. 2020, *PASA*, **37**, e048
- Mohan, N., & Rafferty, D. 2015, *Astrophysics Source Code Library* [record ascl:1502.007]
- Newville, M., Stensitzki, T., Allen, D. B., et al. 2016, *Astrophysics Source Code Library* [record ascl:1606.014]
- Peebles, P. J. E. 1980, *The Large-Scale Structure of the Universe* (Princeton: Princeton University Press)
- Remazeilles, M., Dickinson, C., Banday, A. J., Bigot-Sazy, M.-A., & Ghosh, T. 2015, *MNRAS*, **451**, 4311
- Rengelink, R. B., Tang, Y., de Bruyn, A. G., et al. 1997, *A&AS*, **124**, 259
- Rubart, M., & Schwarz, D. J. 2013, *A&A*, **555**, A117
- Secrest, N. J., von Hausegger, S., Rameez, M., et al. 2021, *ApJ*, **908**, L51
- Secrest, N. J., von Hausegger, S., Rameez, M., Mohayaee, R., & Sarkar, S. 2022, *ApJ*, **937**, L31
- Shimwell, T. W., Tasse, C., Hardcastle, M. J., et al. 2019, *A&A*, **622**, A1
- Siewert, T. M., Hale, C., Bhardwaj, N., et al. 2020, *A&A*, **643**, A100
- Siewert, T. M., Schmidt-Rubart, M., & Schwarz, D. J. 2021, *A&A*, **653**, A9
- Singal, A. K. 2011, *ApJ*, **742**, L23
- Singal, A. K. 2021, *Universe*, **7**, 107
- Taylor, M. B. 2005, *ASP Conf. Ser.*, **347**, 29
- Wagenveld, J. 2023, <https://doi.org/10.5281/zenodo.7962923>
- Wagenveld, J. D., Klöckner, H.-R., Gupta, N., et al. 2023, *A&A*, **673**, A113
- Wright, E. L., Eisenhardt, P. R., Mainzer, A. K., et al. 2010, *AJ*, **140**, 1868
- Zonca, A., Singer, L. P., Lenz, D., et al. 2019, *J. Open Source Softw.*, **4**, 1298

Bibliography

- Aghanim, N. et al. (2020), *Planck 2018 Results - I. Overview and the Cosmological Legacy of Planck*, [A&A 641 A1](#).
- Algera, H. S. B. et al. (2020), *A Multiwavelength Analysis of the Faint Radio Sky (COSMOS-XS): The Nature of the Ultra-faint Radio Population*, [ApJ 903 139](#).
- Ashton, G. et al. (2019), *BILBY: A User-friendly Bayesian Inference Library for Gravitational-wave Astronomy*, [ApJS 241 27](#).
- Baars, J. W. M., R. Genzel, I. I. K. Pauliny-Toth and A. Witzel (1977), *The Absolute Spectrum of Cas A: An Accurate Flux Density Scale and a Set of Secondary Calibrators.*, [A&A 61 99](#).
- Bengaly, C. A. P., T. M. Siewert, D. J. Schwarz and R. Maartens (2019), *Testing the Standard Model of Cosmology with the SKA: The Cosmic Radio Dipole*, [MNRAS 486 1350](#).
- Bengaly, C. A. P., R. Maartens and M. G. Santos (2018), *Probing the Cosmological Principle in the Counts of Radio Galaxies at Different Frequencies*, [JCAP 2018 031](#).
- Bhatnagar, S., T. J. Cornwell, K. Golap and J. M. Uson (2008), *Correcting Direction-Dependent Gains in the Deconvolution of Radio Interferometric Images*, [A&A 487 419](#).
- Blake, C. and J. Wall (2002), *A Velocity Dipole in the Distribution of Radio Galaxies*, [Nature 416 150](#).
- Blum, R. D. et al. (2016), "The DECAM Legacy Survey", *American Astronomical Society Meeting Abstracts #228*, vol. 228, American Astronomical Society Meeting Abstracts 317.01 317.01.
- Bonaldi, A. et al. (2019), *The Tiered Radio Extragalactic Continuum Simulation (T-RECS)*, [MNRAS 482 2](#).
- Bondi, H. and T. Gold (1948), *The Steady-State Theory of the Expanding Universe*, [MNRAS 108 252](#).
- Bridle, A. H. and F. R. Schwab (1999), "Bandwidth and Time-Average Smearing", *Synthesis Imaging in Radio Astronomy II*, ed. by G. B. Taylor, C. L. Carilli and R. A. Perley, vol. 180, Astronomical Society of the Pacific Conference Series 371.
- Briggs, D. S. (1995), *High Fidelity Deconvolution of Moderately Resolved Sources*, PhD thesis.
- Chiu, H.-Y. (1964), *Gravitational Collapse*, [Physics Today 17 21](#).
- Colin, J., R. Mohayaee, M. Rameez and S. Sarkar (2017), *High-Redshift Radio Galaxies and Divergence from the CMB Dipole*, [MNRAS 471 1045](#).
- Colin, J., R. Mohayaee, M. Rameez and S. Sarkar (2019), *Evidence for Anisotropy of Cosmic Acceleration*, [A&A 631 L13](#).
- Colin, J., R. Mohayaee, S. Sarkar and A. Shafieloo (2011), *Probing the Anisotropic Local Universe and beyond with SNe Ia Data*, [MNRAS 414 264](#).
- Condon, J. J. (1984), *Cosmological Evolution of Radio Sources.*, [ApJ 287 461](#).
- Condon, J. J. (1992), *Radio Emission from Normal Galaxies.*, [Annual Review of Astronomy and Astrophysics 30 575](#).

- Condon, J. J. et al. (1998), *The NRAO VLA Sky Survey*, [AJ 115 1693](#).
- Condon, J. J. and S. M. Ransom (2016), *Essential Radio Astronomy*.
- Copernicus, N. (1543), *De Revolutionibus Orbium Coelestium*.
- Cornwell, T. J., K. Golap and S. Bhatnagar (2005), “W Projection: A New Algorithm for Wide Field Imaging with Radio Synthesis Arrays”, *Astronomical Data Analysis Software and Systems XIV*, ed. by P. Shopbell, M. Britton and R. Ebert, vol. 347, Astronomical Society of the Pacific Conference Series 86.
- Cornwell, T. and E. B. Fomalont (1989), “Self-Calibration”, *Synthesis Imaging in Radio Astronomy*, ed. by R. A. Perley, F. R. Schwab and A. H. Bridle, vol. 6, Astronomical Society of the Pacific Conference Series 185.
- Crawford, F. (2009), *Detecting the Cosmic Dipole Anisotropy in Large-Scale Radio Surveys*, [ApJ 692 887](#).
- Dalang, C. and C. Bonvin (2022), *On the Kinematic Cosmic Dipole Tension*, [MNRAS 512 3895](#).
- Dam, L., G. F. Lewis and B. J. Brewer (2023), *Testing the Cosmological Principle with CatWISE Quasars: A Bayesian Analysis of the Number-Count Dipole*, [525 231](#), arXiv: [2212.07733 \[astro-ph.CO\]](#).
- Darling, J. (2022), *The Universe Is Brighter in the Direction of Our Motion: Galaxy Counts and Fluxes Are Consistent with the CMB Dipole*, [ApJL 931 L14](#).
- Deka, P. P. et al. (2023), *The MeerKAT Absorption Line Survey (MALS) Data Release I: Stokes I Image Catalogs at 1-1.4 GHz*, arXiv: [2308.12347 \[astro-ph\]](#).
- Dewdney, P. E., P. J. Hall, R. T. Schilizzi and T. J. L. W. Lazio (2009), *The Square Kilometre Array, Proceedings of the IEEE* [97 1482](#).
- Dey, A. et al. (2019), *Overview of the DESI Legacy Imaging Surveys*, [AJ 157 168](#).
- Domènech, G., R. Mohayaee, S. P. Patil and S. Sarkar (2022), *Galaxy Number-Count Dipole and Superhorizon Fluctuations*, [JCAP 2022 019](#).
- Doppler, C. (1842), *Ueber Das Farbige Licht Der Doppelsterne Und Einiger Anderer Gestirne Des Himmels*.
- Dressler, A. (1991), *The Great Attractor: Do Galaxies Trace the Large-Scale Mass Distribution?*, [Nature 350 391](#).
- Duncan, K. J. (2022), *All-Purpose, All-Sky Photometric Redshifts for the Legacy Imaging Surveys Data Release 8*, [MNRAS 512 3662](#).
- Eddington, A. S. (1913), *On a Formula for Correcting Statistics for the Effects of a Known Error of Observation*, [MNRAS 73 359](#).
- Einstein, A. (1905), *On the Electrodynamics of Moving Bodies*, *Annalen der physik* [17 891](#).
- Einstein, A. (1917), *Kosmologische Betrachtungen Zur Allgemeinen Relativitätstheorie*, *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften* 142.
- Ellis, G. F. R. and J. E. Baldwin (1984), *On the Expected Anisotropy of Radio Source Counts*, [MNRAS 206 377](#).
- Ellis, G. F. R. and A. R. King (1973), *Tilted Homogeneous Cosmological Models*, *Communications in Mathematical Physics* [31 209](#).
- Euler, L. (1748), *Introductio in analysin infinitorum*, apud Marcum-Michaellem Bousquet & socios.
- Fanaroff, B. L. and J. M. Riley (1974), *The Morphology of Extragalactic Radio Sources of High and Low Luminosity*, [MNRAS 167 31P](#).
- Foreman-Mackey, D., D. W. Hogg, D. Lang and J. Goodman (2013), *Emcee: The MCMC Hammer*, [PASP 125 306](#).

-
- Franzen, T. M. O. et al. (2015), *ATLAS – I. Third Release of 1.4 GHz Mosaics and Component Catalogues*, *MNRAS* **453** 4020.
- Górski, K. M. et al. (2005), *HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere*, *ApJ* **622** 759.
- Grobler, T. L., C. D. Nunhokee, O. M. Smirnov, A. J. van Zyl and A. G. de Bruyn (2014), *Calibration Artefacts in Radio Interferometry - I. Ghost Sources in Westerbork Synthesis Radio Telescope Data*, *MNRAS* **439** 4030.
- Guandalin, C., J. Piat, C. Clarkson and R. Maartens (2022), *Theoretical Systematics in Testing the Cosmological Principle with the Kinematic Quasar Dipole*, *arXiv e-prints*, [arXiv:2212.04925](https://arxiv.org/abs/2212.04925), [arXiv: 2212.04925](https://arxiv.org/abs/2212.04925) [astro-ph.CO].
- Gupta, N. et al. (2016), “The MeerKAT Absorption Line Survey (MALS)”, *MeerKAT Science: On the Pathway to the SKA 14* 14, [arXiv: 1708.07371](https://arxiv.org/abs/1708.07371) [astro-ph.GA].
- Gupta, N. et al. (2021), *Blind HI and OH Absorption Line Search: First Results with MALS and uGMRT Processed Using ARTIP*, *ApJ* **907** 11.
- Gupta, N. et al. (2022), *MALS SALT-NOT Survey of MIR-selected Powerful Radio-bright AGN at $0 < z < 3.5$* , *ApJ* **929** 108.
- Hale, C. L. et al. (2019), *Radio Source Extraction with ProFound*, *MNRAS* **487** 3971.
- Hale, C. L. et al. (2023), *MIGHTEE: Deep 1.4 GHz Source Counts and the Sky Temperature Contribution of Star-Forming Galaxies and Active Galactic Nuclei*, *MNRAS* **520** 2668.
- Hale, C. L. et al. (2021), *The Rapid ASKAP Continuum Survey Paper II: First Stokes I Source Catalogue Data Release*, *PASA* **38** e058.
- Hamaker, J. P., J. D. Bregman and R. J. Sault (1996), *Understanding Radio Polarimetry. I. Mathematical Foundations*, *A&AS* **117** 137.
- Haslam, C. G. T., C. J. Salter, H. Stoffel and W. E. Wilson (1982), *A 408-MHZ All-Sky Continuum Survey. II. The Atlas of Contour Maps*, *A&AS* **47** 1.
- Heywood, I. et al. (2020), *VLA Imaging of the XMM-LSS/VIDEO Deep Field at 1–2 GHz*, *MNRAS* **496** 3469.
- Heywood, I., M. J. Jarvis and J. J. Condon (2013), *Sample Variance, Source Clustering and Their Influence on the Counts of Faint Radio Sources*, *MNRAS* **432** 2625.
- Hilton, M. et al. (2021), *The Atacama Cosmology Telescope: A Catalog of >4000 Sunyaev-Zel’dovich Galaxy Clusters*, *ApJS* **253** 3.
- Högbom, J. A. (1974), *Aperture Synthesis with a Non-Regular Distribution of Interferometer Baselines*, *A&AS* **15** 417.
- Holschneider, M., R. Kronland-Martinet, J. Morlet and P. Tchamitchian (1989), “A Real-Time Algorithm for Signal Analysis with the Help of the Wavelet Transform”, *Wavelets*, ed. by J.-M. Combes, A. Grossmann and P. Tchamitchian, Berlin, Heidelberg: Springer Berlin Heidelberg 286, ISBN: 978-3-642-97177-8.
- Hoyle, F. (1948), *A New Model for the Expanding Universe*, *MNRAS* **108** 372.
- Hubble, E. and M. L. Humason (1931), *The Velocity-Distance Relation among Extra-Galactic Nebulae*, *ApJ* **74** 43.
- Intema, H. T., P. Jagannathan, K. P. Mooley and D. A. Frail (2017), *The GMRT 150 MHz All-Sky Radio Survey. First Alternative Data Release TGSS ADRI*, *A&A* **598** A78.
- Jansky, K. (1933), *Electrical Disturbances Apparently of Extraterrestrial Origin*, *Proceedings of the Institute of Radio Engineers* **21** 1387.
- Jonas, J. L. (2016), *The MeerKAT Radio Telescope*, *Proceedings of Science* **25**.

- Lacy, M. et al. (2020), *The Karl G. Jansky Very Large Array Sky Survey (VLASS). Science Case and Survey Design*, *PASP* **132** 035001.
- Lemaître, G. (1927), *Un Univers Homogène de Masse Constante et de Rayon Croissant Rendant Compte de La Vitesse Radiale Des Nébuleuses Extra-Galactiques*, *Annales de la Société Scientifique de Bruxelles* **47** 49.
- LSST Science Collaboration et al. (2009), *LSST Science Book, Version 2.0*, [arXiv e-prints](#).
- Matthews, A. M., J. J. Condon, W. D. Cotton and T. Mauch (2021a), *Cosmic Star Formation History Measured at 1.4 GHz*, *ApJ* **914** 126.
- Matthews, A. M., J. J. Condon, W. D. Cotton and T. Mauch (2021b), *Source Counts Spanning Eight Decades of Flux Density at 1.4 GHz*, *ApJ* **909** 193.
- Mauch, T. et al. (2003), *SUMSS: A Wide-Field Radio Imaging Survey of the Southern Sky - II. The Source Catalogue*, *MNRAS* **342** 1117.
- Mauch, T. et al. (2020), *The 1.28 GHz MeerKAT DEEP2 Image*, *ApJ* **888** 61, [arXiv: 1912.06212](#).
- McConnell, D. et al. (2020), *The Rapid ASKAP Continuum Survey I: Design and First Results*, *PASA* **37**.
- Migkas, K. et al. (2021), *Cosmological Implications of the Anisotropy of Ten Galaxy Cluster Scaling Relations*, *A&A* **649** A151.
- Migkas, K. and T. H. Reiprich (2018), *Anisotropy of the Galaxy Cluster X-ray Luminosity-Temperature Relation*, *A&A* **611** A50.
- Mohan, N. and D. Rafferty (2015), *PyBDSF: Python Blob Detection and Source Finder*, *Astrophysics Source Code Library ascl:1502.007*.
- Mohayaee, R., M. Rameez and S. Sarkar (2021), *The Impact of Peculiar Velocities on Supernova Cosmology*, [arXiv: 2003.10420 \[astro-ph\]](#).
- Nadolny, T., R. Durrer, M. Kunz and H. Padmanabhan (2021), *A New Way to Test the Cosmological Principle: Measuring Our Peculiar Velocity and the Large-Scale Anisotropy Independently*, *JCAP* **2021** 009.
- Newville, M. et al. (2016), *Lmfit: Non-Linear Least-Square Minimization and Curve-Fitting for Python*, *Astrophysics Source Code Library ascl:1606.014*.
- Peebles, P. J. E. (1980), *The Large-Scale Structure of the Universe*.
- Penzias, A. A. and R. W. Wilson (1965), *A Measurement of Excess Antenna Temperature at 4080 Mc/s.*, *ApJ* **142** 419.
- Perley, R. A. and B. J. Butler (2017), *An Accurate Flux Density Scale from 50 MHz to 50 GHz*, *ApJS* **230** 7.
- Perley, R. A., F. R. Schwab and A. H. Bridle (1989), *Synthesis Imaging in Radio Astronomy : A Collection of Lectures from the Third NRAO Synthesis Imaging Summer School*, *Synthesis Imaging in Radio Astronomy* **6**.
- Perlmutter, S. et al. (1999), *Measurements of Ω and Λ from 42 High-Redshift Supernovae*, *ApJ* **517** 565.
- Qin, F., D. Parkinson, C. Howlett and K. Said (2021), *Cosmic Flow Measurement and Mock Sampling Algorithm of Cosmicflows-4 Tully-Fisher Catalog*, *ApJ* **922** 59.
- Rau, U. and T. J. Cornwell (2011), *A Multi-Scale Multi-Frequency Deconvolution Algorithm for Synthesis Imaging in Radio Interferometry*, *A&A* **532** A71.
- Remazeilles, M., C. Dickinson, A. J. Banday, M.-A. Bigot-Sazy and T. Ghosh (2015), *An Improved Source-Subtracted and Destriped 408-MHz All-Sky Map*, *MNRAS* **451** 4311.

-
- Rengelink, R. B. et al. (1997), *The Westerbork Northern Sky Survey (WENSS), I. A 570 Square Degree Mini-Survey around the North Ecliptic Pole*, *A&AS* **124** 259.
- Riess, A. G. et al. (1998), *Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant*, *AJ* **116** 1009.
- Rubart, M. and D. J. Schwarz (2013), *Cosmic Radio Dipole from NVSS and WENSS*, *A&A* **555** 1, arXiv: [1301.5559v2](https://arxiv.org/abs/1301.5559v2).
- Ryle, M. and P. a. G. Scheuer (1955), *The Spatial Distribution and the Nature of Radio Stars*, *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* **230** 448.
- Sault, R. J., J. P. Hamaker and J. D. Bregman (1996), *Understanding Radio Polarimetry. II. Instrumental Calibration of an Interferometer Array*, *A&AS* **117** 149.
- Secrest, N. J., S. von Hausegger, M. Rameez, R. Mohayaee and S. Sarkar (2022), *A Challenge to the Standard Cosmological Model*, *ApJL* **937** L31.
- Secrest, N. J. et al. (2021), *A Test of the Cosmological Principle with Quasars*, *ApJL* **908** L51.
- Shimwell, T. W. et al. (2019), *The LOFAR Two-metre Sky Survey. II. First Data Release*, *A&A* **622** A1.
- Siewert, T. M., M. Schmidt-Rubart and D. J. Schwarz (2021), *Cosmic Radio Dipole: Estimators and Frequency Dependence*, *A&A* **653** A9.
- Siewert, T. M. et al. (2020), *One- and Two-point Source Statistics from the LOFAR Two-metre Sky Survey First Data Release*, *A&A* **643** A100.
- Singal, A. K. (2023), *Discordance of Dipole Asymmetries Seen in Recent Large Radio Surveys with the Cosmological Principle*, *MNRAS* **524** 3636.
- Singal, A. K. (2011), *Large Peculiar Motion of the Solar System from the Dipole Anisotropy in Sky Brightness Due to Distance Radio Sources*, *ApJ* **742** L23.
- Singal, A. K. (2019), *Large Disparity in Cosmic Reference Frames Determined from the Sky Distributions of Radio Sources and the Microwave Background Radiation*, *PhRvD* **100** 063501.
- Singal, A. K. (2021), *Our Peculiar Motion Inferred from Number Counts of Mid Infra Red AGNs and the Discordance Seen with the Cosmological Principle*, *Universe* **7** 107.
- Smolčić, V. et al. (2017), *The VLA-COSMOS 3 GHz Large Project: Continuum Data and Source Catalog Release*, *A&A* **602** A1.
- Stokes, G. G. (1851), *On the Composition and Resolution of Streams of Polarized Light from Different Sources*, *Transactions of the Cambridge Philosophical Society* **9** 399.
- Taylor, A. R., J. M. Stil and C. Sunstrum (2009), *A Rotation Measure Image of the Sky*, *ApJ* **702** 1230.
- Taylor, A. and L. Legodi (2021), *Data Products for Polarization of MeerKAT Calibrators*.
- The CASA Team et al. (2022), *CASA, the Common Astronomy Software Applications for Radio Astronomy*, *PASP* **134** 114501.
- Tiwari, P., S. Ghosh and P. Jain (2019), *The Galaxy Power Spectrum from TGSS ADR1 and the Effect of Flux Calibration Systematics*, *ApJ* **887** 175.
- Tiwari, P. and P. Jain (2013), *Dipole Anisotropy in Integrated Linearly Polarized Flux Density in NVSS Data*, *MNRAS* **447** 2658, arXiv: [1308.3970](https://arxiv.org/abs/1308.3970).
- Tiwari, P., R. Kothari, A. Naskar, S. Nadkarni-Ghosh and P. Jain (2015), *Dipole Anisotropy in Sky Brightness and Source Count Distribution in Radio NVSS Data*, *Astroparticle Physics* **61** 1.
- Tsagas, C. G. and M. I. Kadlitzoglou (2015), *Deceleration Parameter in Tilted Friedmann Universes*, *PhRvD* **92** 043515.
- Urry, C. M. and P. Padovani (1995), *Unified Schemes for Radio-Loud Active Galactic Nuclei*, *PASP* **107** 803.

- van der Vlugt, D. et al. (2021), *An Ultradeep Multiband VLA Survey of the Faint Radio Sky (COSMOS-XS): Source Catalog and Number Counts*, [ApJ 907 5](#).
- Villiers, M. S. de and W. D. Cotton (2022), *MeerKAT Primary-beam Measurements in the L Band*, [AJ 163 135](#).
- Wagenveld, J. D. et al. (2023a), *The MeerKAT Absorption Line Survey: Homogeneous Continuum Catalogues towards a Measurement of the Cosmic Radio Dipole*, [A&A 673 A113](#).
- Wagenveld, J. D., H.-R. Klöckner and D. J. Schwarz (2023b), *The Cosmic Radio Dipole: Bayesian Estimators on New and Old Radio Surveys*, [A&A 675 A72](#).
- Wagenveld, J. (2023), *JonahDW/Bayesian-dipole: First Release*, Zenodo.
- Watkins, R. et al. (2023), *Analysing the Large-Scale Bulk Flow Using Cosmicflows4: Increasing Tension with the Standard Cosmological Model*, [MNRAS 524 1885](#).
- Whittam, I. H. et al. (2022), *MIGHTEE: The Nature of the Radio-Loud AGN Population*, [MNRAS 516 245](#).
- Wilman, R. J. et al. (2008), *A Semi-Empirical Simulation of the Extragalactic Radio Continuum Sky for next Generation Radio Telescopes*, [MNRAS 388 1335](#).
- Wright, E. L. et al. (2010), *The Wide-field Infrared Survey Explorer (Wise): Mission Description and Initial on-Orbit Performance*, [AJ 140 1868](#), arXiv: [1008.0031](#).

Summary

The Copernican principle states that the earth, and therefore humanity, is not in a special place in the Universe. This principle is named after Nicolaus Copernicus, who in 1543 published his heliocentric model, where not the earth but the sun is at the centre of the solar system. With this, the earth was also no longer the centre of the Universe, something that has fundamentally changed our worldview since then. This principle is also the basis of modern cosmology, the branch of astronomy that deals with structure and evolution of the Universe as a whole. This *cosmological principle* states that because we do not occupy a special place in the Universe, the Universe is *isotropic* (it looks the same in all directions) and *homogeneous* (it looks the same in all places). While this is obviously not the case in our immediate environment, we see that this principle holds when we consider the very largest structures in the universe. So, using this principle, we can argue that the measurements we make in cosmology from Earth are representative of the Universe as a whole. While we have good reason to believe that this principle is indeed true, if it were found to be otherwise, it would have significant implications on cosmology.

The cosmic dipole

Because light is not infinitely fast, we look back in time as we look deeper into the Universe. This means that if we look deep enough into the Universe, we can look back to the beginning of the Universe (see also Figure 1). The cosmic microwave background radiation (CMB) is the oldest light from the Universe that we can receive, emitted just 400,000 years after the Big Bang. Before this time, matter in the Universe was so densely packed that light could not escape. The CMB is one of the most studied phenomena in cosmology because it shows a snapshot of the Universe as it was just after the Big Bang. Therefore, it is an important benchmark for cosmological models. However, before a measurement of the CMB can be made, a problem caused by a more local effect must be taken into account. The *doppler effect* describes the change in the perceived wavelength of light or sound what we receive when we move relative to its source, well known from the change in pitch of the sirens of an ambulance driving past. Because our Solar system also moves with respect to the CMB, we see a blueshift in the direction we move in (the wavelength of light becomes shorter), and a redshift in the direction we move away from (the wavelength of light becomes longer). The maximum changes in wavelength are on opposite sides of the sky (see Figure 2). Therefore, we call the observed effect a *dipole* because, like a magnet, it has two opposite poles. To see the finer details in the CMB, one corrects for the velocity of the solar system to remove this cosmic dipole. Remarkably, outside the dipole and the small fluctuations that are so important for cosmological measurements, the background



Figure 1: Schematic map of the Universe as we observe it, with our solar system at the centre, further out the scales get exponentially larger, and we see further back in time. The cosmic microwave background is indicated by the red line and shows us what the Universe looked like just 400,000 years after the Big Bang. The galaxies we use for the measurements in this thesis are largely in the blue band.

radiation looks the same everywhere. This is therefore in good agreement with the cosmological principle, since radiation comes from all sides of the universe.

For a long time, CMB measurements have been the most important and accurate measurements with which to test our cosmological models. Recently, however, more and more precise measurements have been taken from objects closer to us. Many of these measurements match the predictions made with the CMB measurements, but there are also measurements that differ. These measurements can now be made so precisely that these differences can no longer be explained by statistical uncertainties. This suggests that something is missing in our model of the Universe. One of these so-called ‘tensions’ between measurements from the CMB and closer objects is that of the cosmic dipole, and is the subject of this thesis. The motion of our solar system causes a dipole not only in the CMB, but also in other objects. For instance, the Doppler effect causes galaxies to be brighter in the direction we are moving towards, and conversely less bright in the direction we are moving away from. Similarly, one can measure the dipole by counting distant galaxies. However, measurements of this dipole arrive at a velocity many times greater than our measurements with the CMB. This could indicate that the dipole we find in galaxies is not only due to our motion, but that there is also an actual dipole in the number




Figure 2: Map of the cosmic microwave background as seen by the *Planck* satellite. The map is dominated by emission from the Milky Way and by the cosmic dipole, which produces a variation across the sky. The dipole is believed to be caused entirely by the motion of our solar system relative to the cosmic background radiation.

of galaxies in the Universe. This could have drastic implications for cosmology, as a variation on such a huge scale is in direct contradiction to the cosmological principle.

Radio observations

To measure the dipole, it is important that galaxies are far away, which is why we look at galaxies that emit a lot of radio radiation. Radio emission in these galaxies is often associated with an active black hole, which are located in the centres of these galaxies and usually weigh millions to billions of times more than our Sun. These galaxies are so bright that we can pick up radio emission from galaxies billions of light years away. So when we look at radio waves with a wavelength of around 20 centimetres, the brightest sources in the sky are such galaxies. We need more than 200,000 galaxies to measure the small variation in the number of galaxies induced by the dipole. Also, our observations need to cover a large enough part of the sky to see both sides of the dipole. We can achieve this using radio *surveys*. With a radio telescope, by combining many observations over a large part of the sky, a large catalogue of radio sources can be produced. Because of the long wavelengths of radio waves, telescopes are needed with a huge diameter to achieve good resolution. However, we can also achieve this by combining several smaller telescopes in a process we call *interferometry*. It takes many steps to get from interferometric observations to a catalogue of galaxies, and small systematic errors in this process can lead to errors and uncertainties in the measurement of the dipole. Often when such effects



Figure 3: Several individual antennas of MeerKAT, a radio telescope in South Africa. Credit: SARAO.

are identified, they have the strongest impact on the least bright sources. By removing these sources, these effects are thereby also removed. However, this ensures that far fewer sources are available for a measurement of the dipole, resulting in a less precise measurement.

This thesis

In this thesis, I focus on measuring the dipole of galaxies through radio observations, and give special attention to the effects in the data that might affect this measurement. To ensure that the dipole we measure is real and not due to errors introduced during observation, calibration, imaging, or cataloguing the data, I scrutinise this process. This allows me to make a measurement of the dipole with more certainty, and with any luck it also ensures that fewer sources need to be discarded. To make this happen, every step in the process has to be controlled, which requires intimate knowledge of the telescope and survey in question. For this survey, I am working with data from the radio telescope *MeerKAT*, a radio telescope in South Africa that has been making observations since 2018 (see Figure 3). MeerKAT consists of 64 individual antennas, which can be combined with interferometry to form one large radio telescope. The survey I work with is the *MeerKAT Absorption Line Survey* (MALS), which, unlike many other surveys that have been used to make a measurement of the dipole, covers a relatively small fraction of the sky. However, because of the sensitivity of MeerKAT, enough sources are detected to make a measurement of the dipole.

In **Chapter 3**, I examine and describe in detail the process of observations, calibration and creation of images and catalogues of ten individual observations of MALS. There are between 500 and 3,000 sources in each field, depending on how strong the noise is in the observation. Strong noise in observations is often caused by bright sources. Using various techniques, I measure how complete our observations are, or what percentage of sources we find compared to how many there actually are.

Again, this depends on the noise in the observations. Although the observations are less complete for less clear sources, these measurements of catalogue completeness can be used to obtain the correct number of sources. With this, there are enough sources in the complete survey to make a measurement of the dipole.

In **Chapter 4**, I introduce new statistical methods to measure the dipole, specially made to account for the relatively small fraction of the sky that MALS observes. These statistical methods are based on *Bayesian* statistics, which can better deal with incomplete information. Besides a standard method that only takes into account the effect of the dipole, I also introduce a method that takes into account the noise present in the observations, and a method that can combine multiple catalogues to make an even better measurement of the dipole. Before these can be applied to MALS, I test the methods on two other surveys, the *NRAO VLA Sky Survey* (NVSS) and the *Rapid ASKAP Continuum Survey* (RACS). The results are in agreement with previous results, with a dipole three times larger than that in the CMB. For the most significant result, I combine NVSS and RACS, and 800,000 galaxies can be used for a measurement of the dipole. Again, I find the same result, with the highest precision for such a measurement with radio sources so far.

In **Chapter 5**, I combine the knowledge gathered in the previous chapters to make a measurement of the dipole with MALS. Out of a total of 391 observations of MALS, I use 345, the remaining observations have too strong noise or are too close to the disk of the Milky Way. However, I find an effect in the MALS data that causes a variation of the number of detected sources in each observation, thus preventing a measurement of the dipole. The strength of this effect depends on how high above the horizon each field was observed. To apply our statistical method, they must be adjusted to take this effect into account. However, because the effect takes almost the same shape as a dipole, additional assumptions must be made to achieve a measurement of the dipole. Although once again I measure with high precision a dipole many times larger than that of the CMB, the reliability of these measurements is less than the measurements of the previous chapter. However, there are data reprocessing techniques that could reduce the effect, which can be applied to make the measurement more reliable.

In short, we have seen here that even with good preparations, unexpected setbacks can throw a spanner in the works of our measurements. Although the measurement of the dipole with MALS suffered somewhat from the described effect, thanks to the modification on our statistical measurement methods, I was still able to make this measurement. These statistical methods also helped us get the most significant measurement of the dipole with radio galaxies so far. Using these and other measurements, it is now clear that the dipole cannot be easily explained by errors in the data. In all this time, however, it has not become clearer where the difference between the dipole in the CMB and that in radio galaxies comes from. Therefore, this is where future studies on the dipole will have to answer. If this is a cosmological effect, it could have drastic consequences for cosmology, and our place in the Universe.

Zusammenfassung

Das Kopernikus-Prinzip besagt, dass die Erde und damit die Menschheit keinen besonderen Platz im Universum einnimmt. Dieser Grundsatz ist nach Nikolaus Kopernikus benannt, der 1543 sein heliozentrisches Modell veröffentlichte, bei dem nicht die Erde, sondern die Sonne im Mittelpunkt des Sonnensystems steht. Damit war auch die Erde nicht mehr der Mittelpunkt des Universums, was unser Weltbild seither grundlegend verändert hat. Dieses Prinzip ist auch die Grundlage der modernen Kosmologie, dem Zweig der Astronomie, der sich mit dem Aufbau und der Entwicklung des Universums als Ganzes befasst. Dieses *kosmologische Prinzip* besagt, dass das Universum *isotrop* (es sieht in allen Richtungen gleich aus) und *homogen* (es sieht an allen Orten gleich aus) ist, weil wir uns nicht an einem besonderen Ort im Universum befinden. Während dies in unserer unmittelbaren Umgebung natürlich nicht der Fall ist, sehen wir, dass dieses Prinzip gilt, wenn wir die allergrößten Strukturen im Universum betrachten. Mit Hilfe dieses Prinzips können wir also argumentieren, dass die Messungen, die wir in der Kosmologie von der Erde aus vornehmen, repräsentativ für das Universum als Ganzes sind. Wir haben guten Grund zu der Annahme, dass dieses Prinzip tatsächlich zutrifft. Sollte sich jedoch herausstellen, dass dies nicht der Fall ist, hätte dies erhebliche Auswirkungen auf die Kosmologie.

Der kosmische Dipol


Da Licht nicht unendlich schnell ist, blicken wir in der Zeit zurück, wenn wir tiefer in das Universum blicken. Das bedeutet, dass wir, wenn wir tief genug ins Universum schauen, bis zum Anfang des Universums zurückblicken können (siehe auch Figur 1). Die kosmische Mikrowellenhintergrundstrahlung (KMH) ist das älteste Licht aus dem Universum, das wir empfangen können. Es wurde nur 400 000 Jahre nach dem Urknall ausgesandt. Vor dieser Zeit war die Materie im Universum so dicht gepackt, dass das Licht nicht entweichen konnte. Das KMH ist eines der am meisten untersuchten Phänomene in der Kosmologie, weil es eine Momentaufnahme des Universums kurz nach dem Urknall zeigt. Daher ist sie ein wichtiger Maßstab für kosmologische Modelle. Bevor jedoch eine Messung des KMH vorgenommen werden kann, muss ein Problem berücksichtigt werden, das durch einen eher lokalen Effekt verursacht wird. Der Dopplereffekt beschreibt die Änderung der wahrgenommenen Wellenlänge von Licht oder Schall, die wir empfangen, wenn wir uns relativ zu ihrer Quelle bewegen, wie wir es beispielsweise von der Änderung der Tonhöhe der Sirenen eines vorbeifahrenden Krankenwagens kennen. Da sich unser Sonnensystem auch in Bezug auf das KMH bewegt, sehen wir eine Blauverschiebung in der Richtung, in die wir uns bewegen (die Wellenlänge des Lichts wird kürzer), und eine Rotverschiebung in der Richtung, von der wir uns entfernen (die Wellenlänge des Lichts wird länger). Die maximalen Änderungen der Wellenlänge befinden sich auf



Figur 1: Schematische Karte des Universums, wie wir es beobachten, mit unserem Sonnensystem im Zentrum. Weiter draußen werden die Maßstäbe exponentiell größer, und wir sehen weiter zurück in die Vergangenheit. Die kosmische Hintergrundstrahlung ist durch die rote Linie gekennzeichnet und zeigt uns, wie das Universum nur 400 000 Jahre nach dem Urknall aussah. Die Galaxien, die wir für die Messungen in dieser Arbeit verwenden, liegen größtenteils im blauen Band.

entgegengesetzten Seiten des Himmels (siehe Figure 2). Daher nennen wir den beobachteten Effekt einen *dipole*, weil er wie ein Magnet zwei entgegengesetzte Pole hat. Um die feineren Details in der KMH zu sehen, korrigiert man die Geschwindigkeit des Sonnensystems, um diesen kosmischen Dipol zu entfernen. Bemerkenswerterweise sieht die Hintergrundstrahlung außerhalb des Dipols und der kleinen Fluktuationen, die für kosmologische Messungen so wichtig sind, überall gleich aus. Dies steht also in guter Übereinstimmung mit dem kosmologischen Prinzip, da die Strahlung von allen Seiten des Universums kommt.

Lange Zeit waren die KMH-Messungen die wichtigsten und genauesten Messungen, mit denen wir unsere kosmologischen Modelle überprüfen konnten. In letzter Zeit wurden jedoch immer präzisere Messungen an Objekten in unserer Nähe durchgeführt. Viele dieser Messungen stimmen mit den Vorhersagen aus den KMH-Messungen überein, aber es gibt auch Messungen, die davon abweichen. Diese Messungen können inzwischen so genau durchgeführt werden, dass diese Unterschiede nicht mehr durch statistische Unsicherheiten erklärt werden können. Dies deutet darauf hin, dass in unserem Modell des Universums etwas fehlt. Eine dieser so genannten "Spannen" zwischen den Messungen



Figur 2: Karte der kosmischen Mikrowellenhintergrundstrahlung aus der Sicht des *Planck*-Satelliten. Die Karte wird dominiert von der Emission der Milchstraße und dem kosmischen Dipol, der eine Variation über den Himmel erzeugt. Es wird angenommen, dass der Dipol ausschließlich durch die Bewegung unseres Sonnensystems relativ zur kosmischen Hintergrundstrahlung verursacht wird..

der KMH und näheren Objekten ist die des kosmischen Dipols und ist Gegenstand dieser Arbeit. Die Bewegung unseres Sonnensystems verursacht einen Dipol nicht nur in der KMH, sondern auch in anderen Objekten. Der Dopplereffekt bewirkt beispielsweise, dass Galaxien in der Richtung, auf die wir uns zubewegen, heller sind und umgekehrt in der Richtung, von der wir uns entfernen, weniger hell. In ähnlicher Weise kann man den Dipol messen, indem man weit entfernte Galaxien zählt. Die Messungen dieses Dipols erfolgen jedoch mit einer Geschwindigkeit, die um ein Vielfaches höher ist als unsere Messungen mit der KMH. Dies könnte darauf hindeuten, dass der Dipol, den wir in Galaxien finden, nicht nur auf unsere Bewegung zurückzuführen ist, sondern dass es auch einen tatsächlichen Dipol in der Anzahl der Galaxien im Universum gibt. Dies könnte drastische Auswirkungen auf die Kosmologie haben, da eine Variation auf einer so großen Skala in direktem Widerspruch zum kosmologischen Prinzip steht.

Radiobeobachtungen

Für die Messung des Dipols ist es wichtig, dass die Galaxien weit entfernt sind, weshalb wir uns Galaxien ansehen, die viel Radiostrahlung aussenden. Die Radiostrahlung in diesen Galaxien wird oft mit einem aktiven Schwarzen Loch in Verbindung gebracht, das sich im Zentrum dieser Galaxien befindet und in der Regel Millionen bis Milliarden Mal mehr wiegt als unsere Sonne. Diese Galaxien



Figur 3: Einige Einzelantennen des MeerKAT, eines Radioteleskops in Südafrika. Kredit: SARAO.

sind so hell, dass wir die Radiostrahlung von Galaxien auffangen können, die Milliarden von Lichtjahren entfernt sind. Wenn wir also Radiowellen mit einer Wellenlänge von etwa 20 Zentimetern betrachten, sind die hellsten Quellen am Himmel solche Galaxien. Wir brauchen mehr als 200.000 Galaxien, um die kleine Variation in der Anzahl der Galaxien zu messen, die durch den Dipol verursacht wird. Außerdem müssen unsere Beobachtungen einen ausreichend großen Teil des Himmels abdecken, um beide Seiten des Dipols zu sehen. Dies können wir mit Hilfe von Radioteleskopen erreichen. Mit einem Radioteleskop kann durch die Kombination vieler Beobachtungen über einen großen Teil des Himmels ein großer Katalog von Radioquellen erstellt werden. Wegen der langen Wellenlängen der Radiowellen werden Teleskope mit einem großen Durchmesser benötigt, um eine gute Auflösung zu erreichen. Wir können dies jedoch auch durch die Kombination mehrerer kleinerer Teleskope in einem Prozess erreichen, den wir *interferometry* nennen. Von interferometrischen Beobachtungen bis zu einem Galaxienkatalog sind viele Schritte erforderlich, und kleine systematische Fehler in diesem Prozess können zu Fehlern und Unsicherheiten bei der Messung des Dipols führen. Wenn solche Effekte festgestellt werden, wirken sie sich oft am stärksten auf die weniger hellen Quellen aus. Indem man diese Quellen entfernt, werden auch diese Effekte beseitigt. Dies führt jedoch dazu, dass viel weniger Quellen für eine Messung des Dipols zur Verfügung stehen, was zu einer weniger genauen Messung führt.

Diese Dissertation

In dieser Arbeit konzentriere ich mich auf die Messung des Dipols von Galaxien durch Radiobeobachtungen und widme den Effekten in den Daten, die diese Messung beeinflussen könnten, besondere Aufmerksamkeit. Um sicherzustellen, dass der von uns gemessene Dipol real ist und nicht auf Fehler zurückzuführen ist, die bei der Beobachtung, Kalibrierung, Abbildung oder Katalogisierung der Daten

aufzutreten, untersuche ich diesen Prozess. So kann ich den Dipol mit größerer Sicherheit messen und mit etwas Glück auch sicherstellen, dass weniger Quellen verworfen werden müssen. Um dies zu erreichen, muss jeder Schritt des Prozesses kontrolliert werden, was eine genaue Kenntnis des betreffenden Teleskops und der Durchmusterung erfordert. Für diese Durchmusterung arbeite ich mit Daten des Radioteleskops MeerKAT, einem Radioteleskop in Südafrika, das seit 2018 Beobachtungen durchführt (siehe Figur 3). MeerKAT besteht aus 64 Einzelantennen, die mit Interferometrie zu einem großen Radioteleskop kombiniert werden können. Die Durchmusterung, mit der ich arbeite, ist die *MeerKAT Absorption Line Survey* (MALS), die im Gegensatz zu vielen anderen Durchmusterungen, die zur Messung des Dipols verwendet wurden, einen relativ kleinen Teil des Himmels abdeckt. Aufgrund der Empfindlichkeit von MeerKAT werden jedoch genügend Quellen entdeckt, um eine Messung des Dipols vorzunehmen.

In **Kapitel 3** untersuche und beschreibe ich ausführlich den Prozess der Beobachtung, Kalibrierung und Erstellung von Bildern und Katalogen von zehn Einzelbeobachtungen von MALS. In jedem Feld gibt es zwischen 500 und 3.000 Quellen, je nachdem, wie stark das Rauschen in der Beobachtung ist. Starkes Rauschen in Beobachtungen wird oft durch helle Quellen verursacht. Mit verschiedenen Techniken messe ich, wie vollständig unsere Beobachtungen sind, oder wie viel Prozent der Quellen wir im Vergleich zur tatsächlichen Anzahl finden. Dies hängt wiederum vom Rauschen in den Beobachtungen ab. Obwohl die Beobachtungen bei weniger klaren Quellen weniger vollständig sind, können diese Messungen der Vollständigkeit des Katalogs verwendet werden, um die korrekte Anzahl der Quellen zu ermitteln. Damit sind in der vollständigen Durchmusterung genügend Quellen vorhanden, um eine Messung des Dipols vorzunehmen.

In **Kapitel 4** stelle ich neue statistische Methoden zur Messung des Dipols vor, die speziell für den relativ kleinen Teil des Himmels, den MALS beobachtet, entwickelt wurden. Diese statistischen Methoden beruhen auf der *Bayes'schen* Statistik, die besser mit unvollständigen Informationen umgehen kann. Neben einer Standardmethode, die nur den Effekt des Dipols berücksichtigt, stelle ich auch eine Methode vor, die das Rauschen in den Beobachtungen berücksichtigt, sowie eine Methode, die mehrere Kataloge kombinieren kann, um eine noch bessere Messung des Dipols zu erreichen. Bevor diese Methoden auf MALS angewendet werden können, teste ich sie an zwei anderen Durchmusterungen, dem *NRAO VLA Sky Survey* (NVSS) und dem *Rapid ASKAP Continuum Survey* (RACS). Die Ergebnisse stimmen mit früheren Resultaten überein, mit einem Dipol, der dreimal größer ist als der in der KMH. Für das aussagekräftigste Ergebnis kombiniere ich NVSS und RACS, und 800.000 Galaxien können für eine Messung des Dipols verwendet werden. Auch hier finde ich das gleiche Ergebnis, mit der bisher höchsten Genauigkeit für eine solche Messung mit Radioquellen.

In **Kapitel 5** kombiniere ich die in den vorherigen Kapiteln gesammelten Erkenntnisse, um eine Messung des Dipols mit MALS durchzuführen. Von insgesamt 391 MALS-Beobachtungen verwende ich 345, die restlichen Beobachtungen haben zu starkes Rauschen oder sind zu nahe an der Milchstraßenscheibe. Ich finde jedoch einen Effekt in den MALS-Daten, der eine Variation der Anzahl der entdeckten Quellen in jeder Beobachtung verursacht und somit eine Messung des Dipols verhindert. Die Stärke dieses Effekts hängt davon ab, wie hoch über dem Horizont jedes Feld beobachtet wurde. Um unsere statistische Methode anwenden zu können, müssen die Daten angepasst werden, um diesen Effekt zu berücksichtigen. Da der Effekt jedoch fast die gleiche Form wie ein Dipol hat, müssen zusätzliche Annahmen getroffen werden, um eine Messung des Dipols zu erreichen. Obwohl ich auch hier mit hoher Genauigkeit einen Dipol messe, der um ein Vielfaches größer ist als der der KMH, ist die Zuverlässigkeit dieser Messungen geringer als die der Messungen im vorigen Kapitel. Es gibt jedoch Techniken zur Wiederaufbereitung der Daten, die den Effekt verringern und

die Messung zuverlässiger machen können.

Kurz gesagt, wir haben hier gesehen, dass selbst bei guter Vorbereitung unerwartete Rückschläge unseren Messungen einen Strich durch die Rechnung machen können. Obwohl die Messung des Dipols mit MALS etwas unter dem beschriebenen Effekt litt, konnte ich dank der Modifizierung unserer statistischen Messmethoden diese Messung dennoch durchführen. Diese statistischen Methoden verhalfen uns auch zu der bisher aussagekräftigsten Messung des Dipols bei Radiogalaxien. Mit diesen und anderen Messungen ist nun klar, dass der Dipol nicht einfach durch Fehler in den Daten erklärt werden kann. In all dieser Zeit ist jedoch nicht klarer geworden, woher der Unterschied zwischen dem Dipol in der KMH und dem in Radiogalaxien kommt. Daher müssen künftige Studien über den Dipol diese Frage beantworten. Wenn es sich um einen kosmologischen Effekt handelt, könnte dies drastische Folgen für die Kosmologie und unseren Platz im Universum haben.

Samenvatting

Het principe van Copernicus stelt dat de aarde, en daarmee de mensheid, zich niet op een speciale plek in het Universum bevindt. Dit principe is genoemd naar Nicolaus Copernicus, die in 1543 zijn heliocentrisch model publiceerde, waar niet de aarde maar de zon zich in het centrum van het zonnestelsel bevindt. Daarmee was de aarde ook niet meer het centrum van het Universum, iets wat ons wereldbeeld sinds die tijd fundamenteel heeft veranderd. Dit principe vormt ook de grondslag van de moderne kosmologie, de tak van sterrenkunde die zich bezighoudt met structuur en de evolutie van het Universum in haar geheel. Dit *kosmologische principe* stelt dat omdat wij ons niet op een speciale plek in het Universum bevinden, het Universum *isotroop* (het ziet er in alle richtingen hetzelfde uit) en *homogeen* (het ziet er op alle plekken er hetzelfde uit) is. Hoewel in onze directe omgeving dit natuurlijk niet het geval is, zien we dat dit principe geldt als we kijken naar de allergrootste structuren in het heelal. Met dit principe kunnen we dus stellen dat de metingen die we in de kosmologie maken vanuit de aarde representatief zijn voor het Universum als geheel. We hebben dus goede redenen om te geloven dat dit principe inderdaad geldt. Als echter het tegendeel zou worden aangetoond, zou dit aanzienlijke gevolgen hebben op de kosmologie.

De kosmische dipool


Omdat licht niet oneindig snel is, kijken we terug in de tijd als we dieper het heelal inkijken. Dit betekent dat als we diep genoeg in het heelal kijken, we terug kunnen kijken naar het begin van het Universum (zie ook Figuur 1). De *kosmische achtergrondstraling* (KAS) is het oudste licht van het Universum wat we kunnen ontvangen, uitgezonden slechts 400,000 jaar na de Oerknal. Voor deze tijd was de materie in het Universum zo dicht op elkaar gepakt dat licht niet kon ontsnappen. De KAS is een van de meest bestudeerde fenomenen in de kosmologie, omdat het een momentopname laat zien van het Universum zoals het was vlak na de Oerknal. Daarom is het een belangrijk ijkpunt voor kosmologische modellen. Voordat een meting van de KAS echter gemaakt kan worden, moet rekening worden gehouden met een probleem dat wordt veroorzaakt door een meer lokaal effect. Het *dopplereffect* beschrijft de verandering in de waargenomen golflengte van licht of geluid wat we ontvangen als we ten opzicht van de bron ervan bewegen, welbekend van de verandering in toonhoogte van bijvoorbeeld de sirenes van ambulance die langsrijdt. Doordat ons zonnestelsel ook beweegt ten opzicht van de KAS, zien we in de richting waarin we bewegen een blauwverschuiving (de golflengte van licht wordt korter), en in de richting waar we van af bewegen een roodverschuiving (de golflengte van licht wordt langer). De maximale veranderingen in golflengte bevinden zich aan tegenovergestelde kanten van de hemel (zie Figuur 2). Daarom noemen we het waargenomen effect een *dipool*, omdat het net als een magneet twee tegenovergestelde polen heeft. Om de fijnere details in de KAS te



Figuur 1: Schematische kaart van het Universum zoals wij het observeren, met ons zonnestelsel in het midden, verder naar buiten worden de schalen exponentieel groter, en zien we verder terug in de tijd. De kosmische achtergrondstraling is aangegeven met de rode lijn en laat ons zien hoe het Universum eruit zag slechts 400,000 jaar na de Oerknal. De sterrenstelsels die we gebruiken voor de metingen in dit proefschrift bevinden zich grotendeels in de blauwe band.

zien, corrigeert men voor de snelheid van het zonnestelsel om deze kosmische dipool te verwijderen. Bijzonder is dat buiten de dipool en de kleine fluctuaties die zo belangrijk zijn voor kosmologische metingen, de achtergrondstraling er overal hetzelfde uitziet. Dit komt dan ook goed overeen met het kosmologische principe, omdat de straling van alle kanten van het heelal komt.

Lange tijd zijn de metingen van de KAS de belangrijkste en meest accurate metingen geweest waarmee we onze kosmologische modellen konden testen. Recentelijk echter zijn er steeds meer en preciezere metingen uit objecten gehaald die dichterbij staan. Veel van deze metingen komen overeen met de voorspellingen die gemaakt worden met de metingen van de KAS, maar er zijn ook metingen die verschillen. Deze metingen kunnen nu zo precies gemaakt worden, dat deze verschillen niet meer verklaard kunnen worden door statistische onzekerheden. Dit suggereert dat er iets mist in ons model van het Universum. Eén van deze zogeheten ‘spanningen’ tussen de metingen van de KAS en dichterbij gelegen objecten is die van de kosmische dipool, en is het onderwerp van dit proefschrift. De beweging van ons zonnestelsel veroorzaakt niet alleen een dipool in de KAS, maar ook in andere objecten. Zo zorgt het dopplereffect ervoor dat sterrenstelsels helderder zijn in de richting waar we



Figuur 2: Kaart van de kosmische achtergrondstraling, zoals gezien door de *Planck* satelliet. De kaart wordt gedomineerd door emissie van de Melkweg en door de kosmische dipool, die een variatie over de gehele hemel teweegbrengt. Er wordt aangenomen dat de dipool volledig wordt veroorzaakt door de beweging van ons zonnestelsel ten opzichte van de kosmische achtergrondstraling.

naartoe bewegen, en juist minder helder zijn in de richting waar we vanaf bewegen. Zo kan men ook de dipool meten door ver gelegen sterrenstelsels te tellen. Metingen van deze dipool komen echter op een snelheid uit die vele malen groter is dan onze metingen met de KAS. Dit zou kunnen aanduiden dat de dipool die we vinden in sterrenstelsels niet alleen door onze beweging komt, maar dat er ook een daadwerkelijke dipool is in het aantal sterrenstelsels in het heelal. Dit zou drastische gevolgen kunnen hebben voor de kosmologie, want een variatie op zo'n enorme schaal is in directe tegenstrijd met het kosmologische principe.

Radio observaties

Voor het meten van de dipool is het belangrijk dat de sterrenstelsels ver weg staan, en daarom kijken we naar de sterrenstelsels die veel radiostraling uitzenden. Radiostraling in deze sterrenstelsels is vaak geassocieerd met een actief zwart gat, die zich in de centra van deze sterrenstelsels bevinden en meestal miljoenen tot miljarden keren zo veel als onze Zon wegen. Deze sterrenstelsels zijn zo helder dat we radiostraling kunnen opvangen van sterrenstelsels die miljarden lichtjaren ver weg zijn. Als we dus kijken naar radiogolven met een golflengte van rond de 20 centimeter zijn de helderste bronnen in de hemel zulke sterrenstelsels. We hebben meer dan 200,000 sterrenstelsels nodig om de kleine variatie in het aantal sterrenstelsels te meten die de dipool teweegbrengt. Ook moeten onze



Figuur 3: Enkele individuele antennes van MeerKAT, een radio telescoop in Zuid-Afrika. Credit: SARA0.

observaties ook een groot genoeg deel van de hemel beslaan om beide kanten van de dipool te kunnen zien. Dit kunnen we bereiken met behulp van radio *surveys*. Met een radio telescoop kan door vele observaties over een groot deel van de hemel te combineren een grote catalogus aan radiobronnen geproduceerd kan worden. Vanwege de lange golflengtes van radiogolven zijn telescopen nodig met een enorm grote diameter om een goede resolutie te bereiken. Dit kunnen we echter ook bereiken door meerdere kleinere telescopen te combineren in een proces dat we *interferometrie* noemen. Er zijn veel stappen nodig om van interferometrische observaties tot aan een catalogus van sterrenstelsels te komen, en kleine systematische fouten in dit proces kunnen leiden tot fouten en onzekerheden in de meting van de dipool. Vaak als zulke effecten worden geïdentificeerd, hebben ze de sterkste invloed op de minst heldere bronnen. Door deze bronnen te verwijderen zijn daarmee deze effecten ook verwijderd. Dit zorgt er echter voor dat veel minder bronnen beschikbaar zijn voor een meting van de dipool, met als gevolg dat de meting minder precies is.

Dit proefschrift

In dit proefschrift focus ik op het meten van de dipool van sterrenstelsels door middel van radio observaties, en geef ik speciale aandacht aan de effecten in de data die deze meting zouden kunnen beïnvloeden. Om er zeker van te zijn dat de dipool die we meten echt is en niet komt door fouten geïntroduceerd tijdens observatie, calibratie, beeldvorming, of het catalogiseren van de data, neem ik dit proces onder de loep. Dit stelt mij in staat om met meer zekerheid een meting van de dipool te maken, en met enig geluk zorgt dit er ook voor dat er minder bronnen weggegooid hoeven te worden. Om dit voor elkaar te krijgen moet elke stap in het proces gecontroleerd worden, wat intieme kennis vergt van de telescoop en survey in kwestie. Voor dit onderzoek werk ik met data van de radio telescoop *MeerKAT*, een radio telescoop in Zuid-Afrika die sinds 2018 observaties uitvoert (zie Figuur 3).

MeerKAT bestaat uit 64 individuele antennes, die met interferometrie kunnen gecombineerd worden tot één grote radio telescoop. De survey waarmee ik werk is de *MeerKAT Absorption Line Survey* (MALS), die in tegenstelling tot veel andere surveys waarmee een meting van de dipool is gemaakt een relatief kleine fractie van de hemel beslaat. Vanwege de gevoeligheid van MeerKAT worden er echter genoeg bronnen gedetecteerd om een meting van de dipool te maken.

In **Hoofdstuk 3** onderzoek en beschrijf ik in detail het proces van de observaties, calibratie en het maken van beelden en catalogi van tien individuele observaties van MALS. Er bevinden zich tussen de 500 en 3000 bronnen in elk veld, afhankelijk van hoe sterk de ruis is in de observatie. Sterke ruis in observaties wordt vaak veroorzaakt door felle bronnen. Met verschillende technieken meet ik hoe compleet onze observaties zijn, oftewel welk percentage van de bronnen we vinden ten opzichte van hoe veel er daadwerkelijk zijn. Ook dit is afhankelijk van de ruis in de observaties. Hoewel de observaties minder compleet zijn voor minder heldere bronnen, kunnen deze metingen van de compleetheid van de catalogus gebruikt worden om het correcte aantal bronnen te verkrijgen. Hiermee zijn er in de volledige survey genoeg bronnen om een meting van de dipool te maken.

In **Hoofdstuk 4** introduceer ik nieuwe statistische methodes om de dipool te meten, speciaal gemaakt om rekening te houden met de relatief kleine fractie van de hemel die MALS observeert. Deze statistische methodes zijn gebaseerd op *Bayesiaanse* statistiek, dat beter kan omgaan met incomplete informatie. Naast een standaardmethode die alleen rekening houdt met het effect van de dipool, introduceer ik ook een methode die rekening houdt met de ruis aanwezig in de observaties, en een methode die meerdere catalogi kan combineren om zo een nog betere meting van de dipool te maken. Voordat deze kan worden toegepast op MALS, test ik de methodes uit op twee andere surveys, de *NRAO VLA Sky Survey* (NVSS) en de *Rapid ASKAP Continuum Survey* (RACS). De resultaten zijn in overeenkomst met eerdere resultaten, met een dipool die drie maal groter is dan die in de KAS. Voor het meest significante resultaat combineer ik NVSS en RACS, en kunnen 800,000 sterrenstelsels gebruikt worden voor een meting van de dipool. Ook hier vind ik eenzelfde resultaat, met de hoogste precisie voor een dergelijke meting met radio bronnen tot nu toe.

In **Hoofdstuk 5** combineer ik de kennis vergaard in de voorgaande hoofdstukken om een meting van de dipool te maken met MALS. Van de in totaal 391 observaties van MALS gebruiken ik er 345, de overige observaties hebben te sterke ruis of staan te dicht bij de schijf van de Melkweg. Ik vind echter een effect in de MALS data die een variatie van het aantal gedetecteerde bronnen in elke observatie veroorzaakt, en daarmee een meting van de dipool verhindert. De sterkte van dit effect is afhankelijk van hoe hoog boven de horizon elk veld is geobserveerd. Om onze statistische methode toe te passen, moeten deze aangepast worden om met dit effect rekening te houden. Echter omdat het effect vrijwel dezelfde vorm aanneemt als een dipool, moeten extra aannames gemaakt worden om een meting van de dipool te verwezenlijken. Hoewel ik eens te meer met hoge precisie een dipool meet die vele malen groter is dan die van de KAS, is de betrouwbaarheid van deze metingen minder groot dan de metingen van het voorgaande hoofdstuk. Wel zijn er technieken om de data opnieuw te verwerken die het effect zouden kunnen verminderen, die kunnen toegepast worden om de meting betrouwbaarder te maken.

Kortom, we hebben hier gezien dat zelfs met goede voorbereidingen onverwachte tegenslagen roet in het eten kunnen gooien van onze metingen. Hoewel de meting van de dipool met MALS enigszins heeft geleden onder het beschreven effect, heb ik met dank aan de aanpassing op onze statistische meetmethode nog steeds deze meting kunnen uitvoeren. Deze statistische methoden hebben ons ook geholpen aan de meest significante meting van de dipool met radio sterrenstelsels tot nog toe. Met behulp van deze en andere metingen, is het nu duidelijk dat de dipool niet makkelijk kan worden verklaard met fouten in de data. In al deze tijd is het echter niet duidelijker geworden

waar het verschil tussen de dipool in de KAS en die in radio sterrenstelsels vandaan komt. Dit is dan ook waar toekomstige onderzoeken naar de dipool antwoord op moeten gaan geven. Als dit een kosmologisch effect is, zou het drastische gevolgen kunnen hebben voor de kosmologie, en onze plaats in het Universum.