



university of
 groningen

faculty of science
 and engineering

UNIVERSITY OF GRONINGEN

MASTER THESIS

Analysing spectral properties of the FR-II radio galaxy 3C330 using sub-arcsecond LOFAR-VLBI data

Author:

Arno LAFONTAINE (S2990261)

Supervisor:

Prof. Dr. J. P. McKEAN

2nd October 2023

Abstract

Radio galaxies are incredibly interesting and complex sources that have been a topic of study since the beginning of radio astronomy itself. High-resolution observations are required to accurately study their jets, hotspots and lobes in addition to the evolution and interaction of these components with the surrounding environment. For this reason, the current models of radio galaxy physics have been primarily based on observations at high frequencies (≥ 1 GHz), sustaining an existing lack of information on the older plasma in the low-frequency regime (≤ 300 MHz).

The International LOw-Frequency ARray Telescope (LOFAR; ILT) offers a solution by combining large bandwidth (121–166 MHz) sub-arcsecond angular resolution capabilities on the smaller scales with high sensitivity on the large scales.

In this thesis we present new 0.3'' resolution images at 144 MHz of the powerful Fanaroff-Riley Type II radio galaxy 3C330. The data was taken from the second public data release (DR2) of the LOFAR Two-metre Sky Survey (LoTSS) and was calibrated using a combination of the newly developed pipeline for LOFAR Very Long Baseline Interferometry (LOFAR-VLBI) and a manual sub-band self-calibration strategy. By calibrating and using archival VLA data, we also present the first high-resolution spectral index map of the 3C330 lobes between 144 and 8500 MHz.

We observe the presence of multiple secondary hotspots that exhibit spectral flattening and by analysis of their spectral index values we find that both the 'dentist-drill' and the precessing jet termination model best fit the spectral evolution of 3C330. Additionally, our results show evidence of a change in precession angle of about 2° , which is smaller than the inferred 6° from other 3C sources.

Contents

Abstract	1
1 Introduction	4
2 Overview of 3C330	5
2.1 Morphology and Spectral Index	7
2.2 Resolved observations of 3C330	8
3 Theory	10
3.1 Synchrotron radiation	10
3.2 LOFAR antennas	11
3.2.1 Beamforming	11
3.3 Radio interferometry	11
3.3.1 Imaging and (de)convolution	15
3.4 System performance and limitations	16
3.4.1 Time and bandwidth smearing	16
3.5 Calibration	16
3.5.1 Direction-dependent effects (DDEs)	17
3.5.2 Instrumental effects (DIEs)	17
4 Observation and data processing	18
4.1 Observation details	18
4.2 Obtaining the data	18
4.3 Data processing requirements	19
4.4 PREFACTOR pipeline	20
4.5 LOFAR-VLBI pipeline	22
4.6 8.5 GHz VLA Observation	25
5 Results	27
5.1 Spectral energy distribution (SED)	27
5.2 Data inspection and flagging	27
5.3 Self-calibrating individual subbands	28
5.3.1 Flux calibration	29
5.4 Creating a concatenated (MT-)MFS model	30
5.4.1 Calibrating all subbands	30
5.5 Spectral index map (144 MHz–8.5 GHz)	34
6 Discussion	36
6.1 Surface brightness maps	36
6.2 Spectral index map (144 MHz – 8.5 GHz)	36
6.2.1 Spectral structure	36
6.2.2 Jet acceleration/termination models	37
6.3 Issues during post-pipeline cleaning and calibration	37
7 Conclusion	38
7.1 Technical issues and future recommendations	38
8 Acknowledgements	39
Appendices	42
A Data inspection	42
A.1 Diagnostic plots of Pre-Facet-Calibrator.parset	42
A.2 Diagnostic plots of Pre-Facet-Target.parset	42
A.3 Diagnostic plots of Delay-Calibration.parset	43
A.4 Diagnostic plots in using CASA's plotms	44

B Complementary results	44
B.1 Sources for the spectral energy distribution (SED)	44
B.2 8.5 GHz VLA observation	45

1 Introduction

The study of radio galaxies is one of the oldest fields in radio astronomy, dating back to the discovery (Grote Reber, 1939) and identification (Baade & Minkowski 1954) of the first radio galaxy Cygnus A. Subsequent observations of other radio galaxies showed that these sources share some similar features that are distinctly recognizable: large, often symmetrical, areas of radio emission around an actively emitting black hole. Additionally, the radio emission of the source is generally much larger than its optical counterpart. The bright emission from radio galaxies originates primarily from the lobes and the hotspots and is a result of shocked electrons precessing about the magnetic field lines inside the cocoon, making them radiate away their energy in the form of photons. Because the luminosity of a synchrotron electron is proportional to its energy, the hotspots will contain younger and therefore brighter electrons, while the lobes typically have a lower surface brightness due to the presence of older electrons.

A first qualitative understanding of radio galaxy morphologies came from a ‘twin-exhaust’ model (Blandford & Rees 1974; Scheuer 1974), in which an active galactic nucleus (AGN) fuels a pair of bipolar relativistic jets that transfer energy to two respective hotspots. The hotspots represent the collision of the particles with the surrounding intergalactic medium (IGM) and, as such, are extremely bright and compact. The highly energetic plasma streams inbetween the nucleus and the hotspots start to slow down when nearing the hotspots, as here the ambient density increases and angular momentum should be conserved. The decelerated material starts to form an expanding lobe or cocoon back towards the nucleus, which, if it reaches supersonic speeds, can impart a shock back into the surrounding medium. As such, the observable components of the classical radio galaxy model include an AGN *core*, two relativistic symmetrical *jets* facing outwards from the magnetic field poles, two *hotspots* at the intersection of the plasma stream with the IGM and two expanding *lobes* that consist of material flowing back towards the nucleus.

This ‘beam’ model was originally based on observations of low-redshift sources and can be considered fairly simplistic as they assumed uniform environments. Nevertheless it remains a valid description of the general underlying mechanisms that form radio galaxies. Subsequent (semi-analytical) models like Turner & Shabala (2015) and Hardcastle (2018) have expanded on it by taking into account properties of the radio environment (e.g. pressure and mass density profiles; see Figure 1).

Besides the refinement of the (semi-)analytical models, several numerical simulations of radio jets were also published in the last decade. These make use of the fact that the relative scale difference between the Larmor radii of jet particles and the scales of the surrounding systems allow for the modelling of AGN jets as fluids or plasmas (Blandford & Rees 1974). The simulation of relativistic jets allows for the addition of detailed physics like e.g. magnetic fields, spectral aging (English et al. 2016; Yates-Jones et al.

2022) or it can be used to study the interaction of relativistic jets with e.g. stars, gas clouds or the IGM (Perucho & López-Miralles 2023).

While current models are constantly improving, there is still a large discrepancy with the observational maps. In the decades since their discovery, and with the advent of higher sensitivity and higher resolution telescope observations (e.g. VLA, e-MERLIN), it had become clear that the evolution dynamics and morphology of radio galaxies can be quite complex.

Many radio sources were first discovered during the Cambridge Radio Surveys and the subsequent releases of the Cambridge Catalogues of Radio Sources, commonly abbreviated to 1C, 2C, 3C, etc (e.g. Edge et al. 1959; Bennett 1962; Laing & Peacock 1980). The new discoveries in these surveys eventually lead to more refined studies of high-resolution maps at 8.4 GHz ($z < 0.15$, Leahy et al. 1997; $0.15 < z < 0.30$, Hardcastle et al. 1997) and studies by e.g. Fernini et al. (1997) that challenged the classical models by comparing radio galaxy structures at 5 GHz to quasars (QSR).

Fanaroff & Riley (1974) were the first to make a distinction between two types of commonly observed radio galaxies, accordingly named Fanaroff-Riley type-I (FR-I) and Fanaroff-Riley type-II (FR-II). This classification scheme takes into account the distance between the brightest components on the lobes and divides that by the total distance to the nearest outer contour on either side. If the calculated ratio is smaller than 0.5 the source is classified as FR-I and if the ratio is larger than 0.5 the source is classified as FR-II. Due to the nature of this classification scheme it can quickly be seen that FR-I sources are brighter near the centers of their respective galaxies and their lobes become fainter as the distance to the core increases. These sources also show steeper spectra, indicating that the radiation coming from FR-I sources has aged significantly. FR-II sources (see Figure 1) are characterized by their high relative brightness at the end of their lobes and by a much flatter spectra at the hotspots. While this is the general description, the FR-II category can also include galaxies with e.g. bended jets, multiple hotspot features, additional jets and compact knots along the lobe.

The characteristic of multiple bright components is one that is commonly seen in a lot of radio galaxies (see e.g. Cygnus A; Perley et al. 1984, 3C351; Laing 1981). The phenomenon can be proposedly explained by the ‘dentist drill’ model of Scheuer (1982), which has been widely accepted and supported by more recent hotspot studies (e.g. Steenbrugge et al. 2008, Pyrzas et al. 2015). The model theorises that either the precession of the jet or a deflection event can alter the direction of the jet, and consequently the location of the primary hotspot. Over time this causes the jet to penetrate further and further into the surrounding medium, leaving behind spots of material that cool by expansion at the locations of the previously fueled relic hotspots.

Due to the limitations in angular resolution of the past instruments, there has been an abundance of detailed spectral information at higher frequencies (> 1 GHz). How-




Figure 1: Schematic overview of the components of a FR-II jet model [Turner & Shabala \(2015\)](#).

ever, currently there still exists a shortage of empirical sub-arcsecond data of the lower frequency regime (< 300 MHz). The nature of synchrotron electron radiation dictates that the lower frequency regime corresponds to the oldest parts in the radio lobes, and thus a combination of high- and low-frequency studies will provide the most information on a radio galaxy’s evolution. Besides this, observations at low-frequencies can make use of the shorter baselines to track old lobe radiation at the larger angular scales. With this information we can test the accuracy of the standard radio galaxy models, as they predict the oldest radiation to reside in the lobe regions furthest away from the hotspot and nearest the core.

Telescopes like the LOw Frequency Array (LOFAR; [van Haarlem et al. 2013a](#)) have been specifically designed for such purposes, combining a high sensitivity ($\sim 100 \mu\text{Jy}/\text{beam}$) with high resolutions of up to $5''$ in the 10–240 MHz frequency range. This telescope has been used to e.g. study the spectral properties of nearby radio galaxies, like the turnover in the spectrum of Cygnus A ($z = 0.0561$; [McKean et al. 2016](#)) and has overall been providing a tremendous amount of data by surveying the northern sky in the LOFAR Two-metre Sky Survey (LoTSS; [Shimwell et al. 2017](#); [Shimwell et al. 2019](#)).

Observing the low-frequency regime with sub-arcsecond angular resolution is however more challenging (especially for more distant sources), due to the inversely proportional relationship between resolution and frequency. A solution comes in the form of large interferometric telescopes that overcome this problem by observing with extreme distances between their antennas, also called Very Long Baseline Interferometry (VLBI). One such telescope is the International LOFAR Telescope (ILT), which makes use of the LOFAR array in the Netherlands plus an additional 13 international stations located across Europe. This large geometric spread allows baselines from the scales of about 0.25 to 1980 km, meaning that for typical low-frequency observations of 150 MHz, the ILT is sensitive to angular scales of about $27.5'$ to $0.21''$.

One of the main issues with the ILT is that the observed interferometric data is quite technically challenging to calibrate. This has however recently been remedied by a new calibration strategy presented by [Morabito et al. \(2022\)](#), which consequently provides a standardized way to obtain

high-resolution calibrated data of many radio galaxies in the northern sky. Examples of studies having successfully used the strategy include the investigation of the evolution of FR-I/FR-II galaxy 3C293 ([Kukreti et al. 2022](#)), testing models for the ring structure in Hercules A ([Timmerman, R. et al. 2022](#)) and environmental studies of powerful radio galaxies 3C34 and 3C320 ([Mahatma et al. 2023](#)). The preliminary success of this strategy shows that it can be used to further explore the spectral context of radio galaxies and give clues about their history and evolution.

In this thesis we present new observations of the large and bright radio galaxy 3C330 at $\sim 0.3''$ resolution taken with the ILT at frequencies from 121 to 166 MHz. We provide the calibrated data at 144 MHz using a combination of the LOFAR-VLBI strategy and manual calibration and also include a derived spectral index map using archival VLA data at 8.5 GHz. Besides a discussion on 3C330’s morphology and a comparison with the standard galaxy evolution models, we also aim to understand and test the best methods of processing an ILT dataset using partially manual calibration methods, specifically for sources with a large angular size ($> 30''$).

Section 2 gives an overview of the source 3C330 including any relevant previous observations. In Section 3 we give an overview of the necessary theory that is needed to understand our results. In Section 4 we describe the new ILT observation of 3C330 and we explain the data processing steps from both the LOFAR-VLBI pipeline ([Morabito et al. 2022](#)) and our manual work outside of the pipeline. Then, in Section 5, we present our results, in Section 6 we discuss the obtained results and finally in Section 7 we summarize our conclusions and give recommendations for similar further studies.

Throughout this thesis, we follow the standard nomenclature for the naming of the lobes and hotspots by [Gilbert et al. \(2004\)](#) and define the spectral index as $S_\nu \propto \nu^\alpha$. At the redshift of 3C330 ($z = 0.54725$), 1 arcsecond corresponds to a projected distance of 6.583 kpc for a $\Omega_M = 0.308$, $\Omega_\Lambda = 0.692$ flat-cosmology with $H_0 = 67.8 \text{ km s}^{-1} \text{ Mpc}^{-1}$.

2 Overview of 3C330

The source 3C330 is a radio galaxy in the northern sky at a moderately high redshift of 0.5490. The nucleus is positioned at a right ascension of $16^{\text{h}}09^{\text{m}}36.607^{\text{s}}$ and a declination of $+65^\circ56'43.61''$ (J2000) and the object is speculated to lie in a group or cluster ([Hardcastle et al. 2002](#)). The object has a relatively weak core (0.7 mJy at 5 GHz; [Ferrini et al. 1997](#)), but has quite powerful radiating lobes with a full-source observed flux density of 30.3 Jy at 178 MHz ([Kellermann et al. 1969](#)).

While 3C330 has been observed many times (mostly as part of large-area surveys), detailed studies on its spectral properties and evolution are quite limited to non-existent. Despite this, the source displays quite some interesting and common characteristics on both the small and large scales that can help with creating a fundamental under-

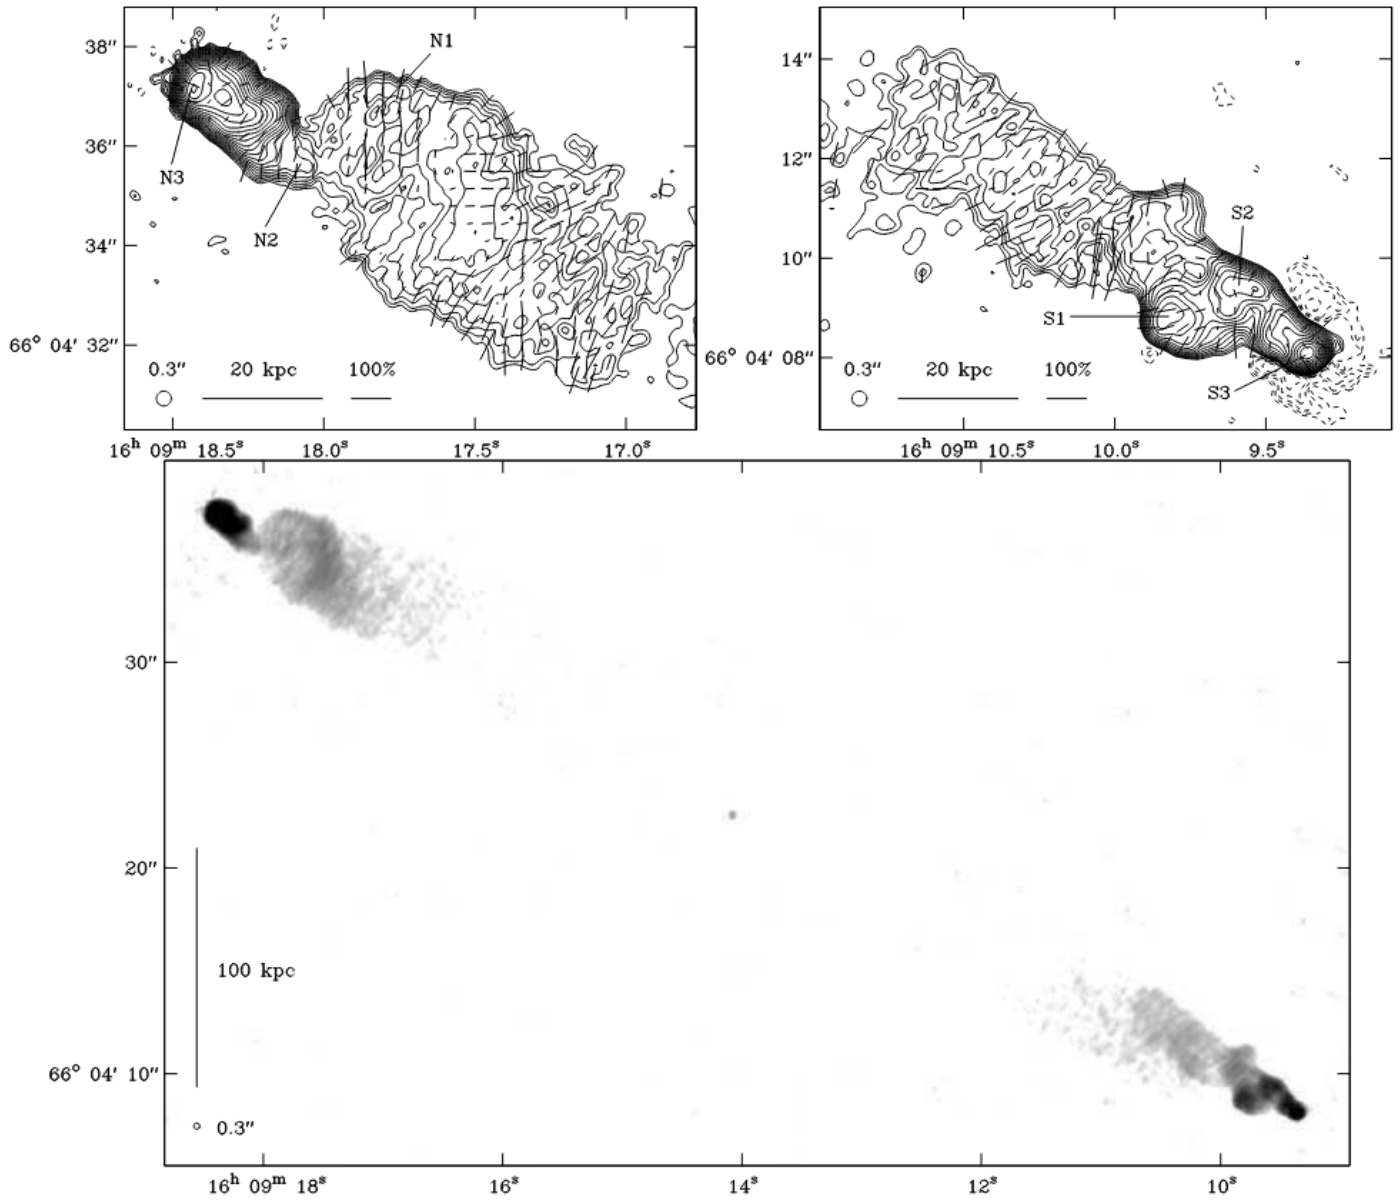


Figure 2: 3C330 observed with the VLA at 8.5 GHz with a resolution of 0.3". Above left shows the northern hotspot, above right shows the southern hotspot and the bottom shows the entire source (Gilbert et al. 2004).

standing of radio galaxies in the low-frequency regime. Due to its relatively large size and brightness the source is also relatively easy to observe, making it a good candidate for testing the ILT capabilities and giving it potential to become an excellent representation for similar more distant radio galaxies ($z \approx 0.5$). The combination of all these factors make 3C330 an ideal target for this study.

Discovery and naming: Together with 470 other radio sources, the object 3C330 was first observed and listed in the Third Cambridge Catalogue of Radio Sources (3C) by [Edge et al. \(1959\)](#). The observations for this survey were made with the Cambridge Interferometer at a frequency of 159 MHz. The 3C catalogue was later revised by [Bennett \(1962\)](#) to the 3CR catalogue based on new observations done at 178 MHz and subsequently even later the 3CR catalogue was revised to 3CRR by [Laing et al. \(1983\)](#) with the goal to increase its completeness and correct for problems with confusion that were still present in the previous two catalogues. Currently the 3CRR catalogue includes all extragalactic radio sources on the northern hemisphere that have a flux density $S_{178\text{MHz}} \geq 10\text{Jy}$, a declination $\delta \geq 10^\circ$ and a galactic latitude $|b| \geq 10^\circ$. This makes the catalogue a reliable sample of all radio galaxies and radio loud quasars present in the northern hemisphere.

The name 3C330, like all object names in the 3CRR catalogue (and all other Cambridge catalogues for that matter) comes from a combination of its catalogue name and its number in order of increasing right ascension in said catalogue. While 3C330 is the common name of this object, its formal or International Astronomical Union (IAU) name is 1609+660, referring to position in right ascension and declination on the sky.

2.1 Morphology and Spectral Index

Morphology: The morphology of 3C330 can be described as a typical FR-II radio galaxy (see Figures 2, 3 and 4). It has two symmetrical hotspots with lobes behind them that taper back towards the galactic nucleus. In generally all resolved observations the northern lobe appears brighter than the southern lobe. Both hotspots lie at a similar distance of about $30''$ from the core, which in turn lies centrally spaced inbetween. With all source components lying along the same axis, 3C330 extends to about $62''$ in length. The width of 3C330 is narrowest at the hotspots ($\sim 5''$), but widens up to about $15''$ at the position of the core.

From sensitive low-resolution observations there is evidence of a constant surface brightness bridge between the two hotspots (e.g. [Leahy et al. 1989](#), [Hardcastle et al. 2002](#)).

From high-resolution 8.5 GHz VLA observations ([Fernini et al. 1997](#), [Fernini 2002](#), [Gilbert et al. 2004](#)) it was shown that in the radio both hotspots can be considered as complex, showing signs of multiple interesting features (see e.g. Figure 2). The northern hotspot is surrounded by an elongated lobe of about $2''$ that expands into a wider region of higher surface brightness, while the southern lobe has a peculiar shape with multiple compact regions of

high surface brightness. For both northern and southern lobe the brightest emission (i.e. primary hotspot) can be found at the furthest end from the core. Still, these high-resolution maps show evidence of potential multiple hotspots, filaments, jet knots and an overall evolving emission structure. Both lobes also show signs of S-like curvature and sharp bends, potentially indicating a changing jet direction.

Spectral Index: With the high number of flux measurements available the global spectral index value for 3C330 has been calculated many times. For the low-frequency radio regime we find values of $\alpha_{750}^{38} = -0.68 \pm 0.01$ ([Kellermann et al. 1969](#)), while the spectral index tends to steepen somewhat at higher frequencies $\alpha_{5000}^{750} = -0.84 \pm 0.03$ ([Laing & Peacock 1980](#)).

More difficult to find are the spectral index properties of the individual source components. The only known spectral index map of 3C330 is the low-resolution $3''$ two-point spectral index map between 151 MHz (MERLIN) and 1.5 GHz (VLA) provided by [Leahy et al. \(1989\)](#). This map can be seen in Figure 3. It shows that the spectral index is highest, or steepest, in the center ($\alpha_{1418}^{151} \approx -1.4$) and starts to flatten when moving outwards towards the emission in the hotspots ($\alpha_{1418}^{151} \approx -0.6$). For a typical FR-II source this behaviour is as expected as it signals the natural decrease in electron age when moving from the older back flow plasma to the newer hotspot emission.

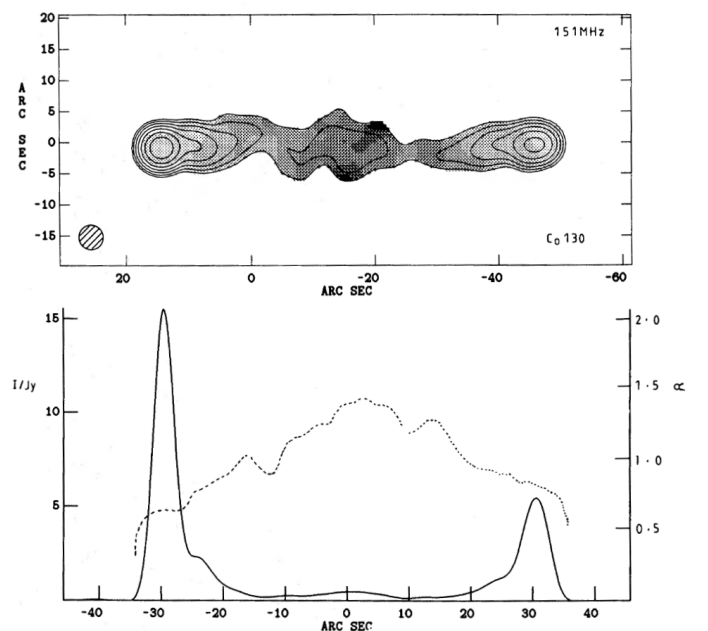


Figure 3: Above: two-point 151–1481 MHz spectral index map of 3C330. The map has been rotated by about 30° such that the source is in line with the axis of the plot. The greyscale colors represent spectral index values from 0.6 (white) to 2.0 (black) in increments of 0.2. Below: profiles of the intensity $I_{151\text{MHz}}$ (solid line) and the spectral index α (dashed line) along a straight line through both hotspots ([Leahy et al. 1989](#)).




Figure 4: Contours show the X-ray emission of 3C330 observed with Chandra. The contours are smoothed by a $6''$ FWHM Gaussian (see lower left corner) and are logarithmically increasing by $\sqrt{2}$. Overlaying in black is 8.4 GHz data from the VLA (Hardcastle et al., 2002).

2.2 Resolved observations of 3C330

As we will be analysing a high angular resolution observation of 3C330, it is useful to study similar observations already available in the literature. There have been multiple resolved observations of 3C330 at several frequencies in (mostly) the radio regime. We describe the most relevant ones here:

X-ray (Chandra): The first X-ray observations of 3C330 were published by Hardcastle et al. (2002). The object was observed with the Chandra X-ray observatory together with two other strong FR-II sources (the quasars 3C 263 and 3C 351). The main goal of this observation was to model the synchrotron spectra of the hotspots to better understand the origin of the emission from these regions. Assuming the dominant source of emission in the hotspots to be synchrotron self-Compton (SSC), the corresponding observed X-ray flux density can be used to estimate a magnetic field in these areas. Observations from the radio-loud FR-II object Cygnus-A indicate that the X-ray emission from its hotspots is actually consistent with the SSC emission when the magnetic field strength is close to the minimum energy or equipartition values. Meanwhile however, observations of other sources indicate that the X-ray emission is too bright to correspond with SSC emission from an equipartition magnetic field. Hardcastle et al. (2002) therefore propose a classification difference between SSC and non-SSC emission objects. The object 3C330 falls in the former category, meaning its X-ray flux density is consistent with SSC emission in the case of a magnetic field close to equipartition. The observed image of 3C330 in the X-ray can be seen in Figure 4.

1.4 GHz (VLA): McCarthy et al. (1991) presents a resolved $3.75''$ resolution image of 3C330 from their VLA

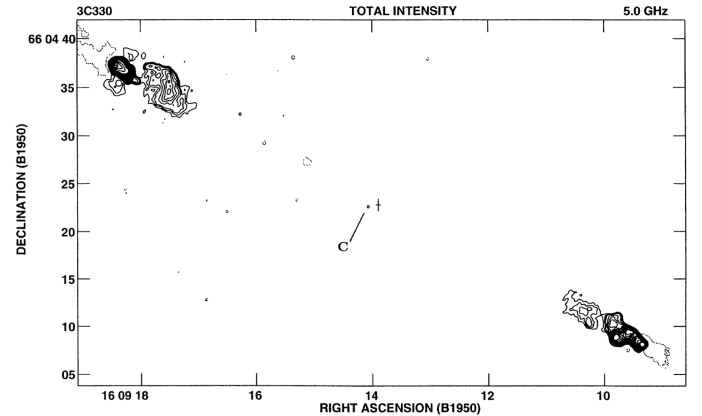


Figure 5: Total intensity $0.43''$ resolution contour map of 3C330 at 5 GHz from combined VLA A and B configuration data. Both lobes are resolved and the core is denoted by the letter 'C' (Fernini et al. 1997).

observation at 1440 MHz. They compare their data with optical $[\text{O II}] \lambda 5007$ emission-line data from Hubble and find that the spectral line emission is asymmetric and concentrated around the North-eastern lobe.

5 GHz (VLA): An observation of 3C330 at 5 GHz using the combined A and B configuration of the VLA is presented by Fernini et al. 1997 (see Figure 5). The map has an angular resolution of 0.43arcsec and shows clearly resolved lobes and hotspots. Additionally, a first detection of a weak compact central feature was made. The observation was part of a study that compared observations of classical FR-II radio galaxies with observations of quasars (QSRs) to test the unification model proposed by Barthel (1989).

8.5 GHz (VLA): An 8.5 GHz VLA observation of 3C330 is presented by Fernini (2002), and later included in a sample of FR-II radio galaxies by Gilbert et al. (2004). The main objective of these studies was to investigate and map the hotspot and lobe structures of such sources using high angular resolution data. Similar to the 5 GHz observations, clear detections of both lobes and hotspots were observed, but now at an angular resolution of about 0.25arcsec . The final map of this observation can be seen in Figure 2.

144 MHz (LOFAR): In preparation for future LOFAR(-VLBI) observations of the northern sky, the LOFAR Long-Baseline Calibrator Survey (LBCS; Jackson et al. 2022) was initiated in 2016. Using LoTSS data, the LBCS aims to create a database with suitable sources to calibrate the highest resolution observations made with the ILT. Such sources should have correlated flux densities at frequencies around 110–190 MHz on scales of a few hundred milliarcseconds (Jackson et al., 2016). The quality of potential calibrators was assessed by the separation distance, detection of the source by the international stations, detection of fringes, signal-to-noise ratio) and the quality of the observation (percentage of detections on the short

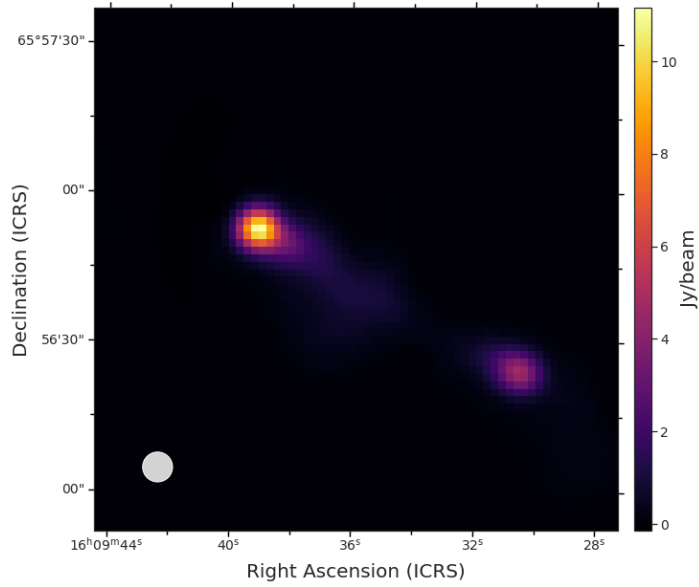


Figure 6: Cutout of 3C330 from DR2 of the LoTSS. The map has been convolved with a circular Gaussian beam of $6''$ resolution.

baselines). From the LBCS database it was found that the source 3C330 agrees with all of these criteria, and is thus one of the chosen calibrator candidates for calibrating LOFAR-VLBI observations.

From the calibrated second data release (DR2) of LoTSS (Shimwell et al. 2022) we can extract an observation of 3C330 with essentially the Dutch array of the ILT ($\sim 6''$) at 144 MHz. The result of this observation can be seen in Figure 6. There is a clear detection of the two lobes, and inbetween there is a fine taper towards some fainter flux at larger angular scales.

3 Theory

In this section we will describe the underlying theory needed to understand our methods and results. This entails a description of synchrotron radiation, an overview of the LOFAR telescope, a general description of radio interferometry and imaging, and finally an explanation of the theory behind the calibration strategy of LOFAR-VLBI.

3.1 Synchrotron radiation

To understand and analyze what we see when we look at low-frequency radio sources, it is important to have some knowledge about the various physical mechanisms that can produce radio emission. Identifying these relevant emission processes gives vital information about the origin, morphology and the ongoing processes in the source. To remain relevant, we will specifically look at the most dominant emission process for observing our source with LOFAR.

Synchrotron radiation is a type of emission that is most dominant in the low-frequency regime. It occurs when charged particles (protons and electrons) of relativistic speeds are captured in a magnetic field and start spiraling around the magnetic field lines. This rotation implies an acceleration perpendicular to the direction of motion. As a direct consequence of Maxwell's equations this means the particle will start emitting radiation. The average isotropic synchrotron power emitted by a relativistic electron over its lifetime can be described by (Rybicki & Lightman, 1979):

$$\langle P_{\text{synch}} \rangle = \frac{4}{3} \sigma_T \beta^2 \gamma^2 B^2 c U_B, \quad (3.1)$$

where $\sigma_T \equiv \frac{8\pi}{3} \left(\frac{e^2}{m_e c^2}\right)^2$ is the Thompson cross-section, $U_B \equiv \frac{B^2}{8\pi}$ is the magnetic energy density, $\beta \equiv v/c$, $\gamma \equiv (1 - \beta^2)^{-1/2}$ is the Lorentz factor and B is the the magnetic field. It is thus important to note that the emitted power is heavily dependent on both the particle velocity v (through β and γ), and the strength of the magnetic field B .

An average electron in a radio source will have an emitting lifetime in the order of 10^3 to 10^6 years before it starts to lose its energy due to synchrotron emission or other processes. Assuming the primary energy loss is due to synchrotron radiation, a first order approximation of the *synchrotron lifetime* can be obtained by the ratio of its total energy ($E_e = \gamma m_e c^2$) by the power it is emitting:

$$\tau \approx 16.4 \text{ years} \times \frac{1}{\gamma} \left[\frac{1 \text{ Gauss}}{B^2} \right]. \quad (3.2)$$

This tells us that electrons with higher energies or those in the presence of stronger magnetic fields lose their energies faster.

The spectrum of synchrotron radiation is a superposition of the synchrotron radiation of all individual emitting electrons. The electron population can be explained

by a power-law distribution $N(E)dE \propto E^{-k}dE$, which consequently means the spectrum also follows a power-law. Here we define $N(E)$ as the number of electrons per unit volume, with energies E to $E + dE$ and k as an arbitrary value representing the power-law energy slope.

We find the *spectral energy distribution (SED)* of a source by plotting its flux density S_ν as a function of frequency ν . When we consider the low-frequency regime of a synchrotron-dominating source, the observed SED can be described by $S_\nu \propto \nu^\alpha$, where the implicit slope is called the *spectral index*:

$$\alpha(\nu) = \frac{d \log S_\nu(\nu)}{d \log \nu}. \quad (3.3)$$

Typical values of the synchrotron spectral index are fairly constant and lie between -0.5 to -0.8 . However, there are two places at both edges where the spectrum deviates from this power-law behaviour.

A low-frequency steepening can be observed due to *synchrotron self-absorption*. Here, at longer wavelengths, the absorption cross-section of electrons becomes large enough that the emitted synchrotron radiation can directly be absorbed again by another electron of the same frequency. The plasma thus becomes optically thick and we are unable to observe beyond the thin shell at the surface. This region has a typical spectral index of $\alpha = 2.5$.

At high frequencies we also observe a sudden steepening of the spectral index to around $\alpha = -1.2$. This is a consequence of higher energy electrons radiating away their energy first, such that they are no longer present. Over time, the frequency at which this cutoff occurs will slowly shift towards lower and lower frequencies. An illustration of a typical radio spectrum with both breaks is illustrated in Figure 7.

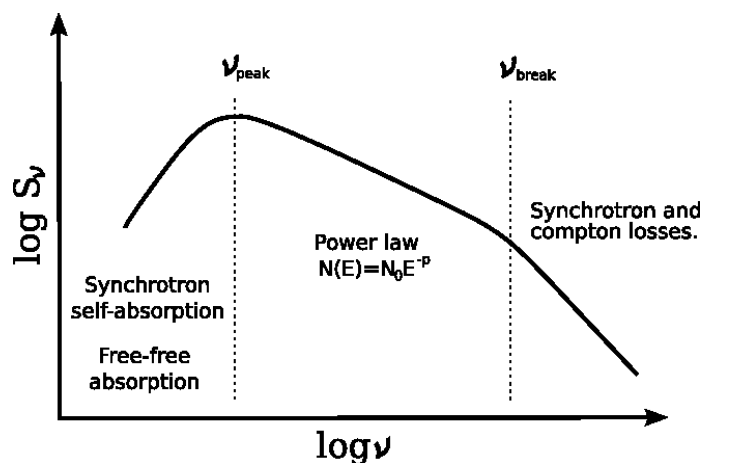


Figure 7: A log-log schematic of a synchrotron spectrum with flux density as a function of frequency. The break at ν_{peak} indicates the transition from optically thick to optically thin, while the break at ν_{break} relates to the age of the emitting electrons (Montenegro-Montes 2009).

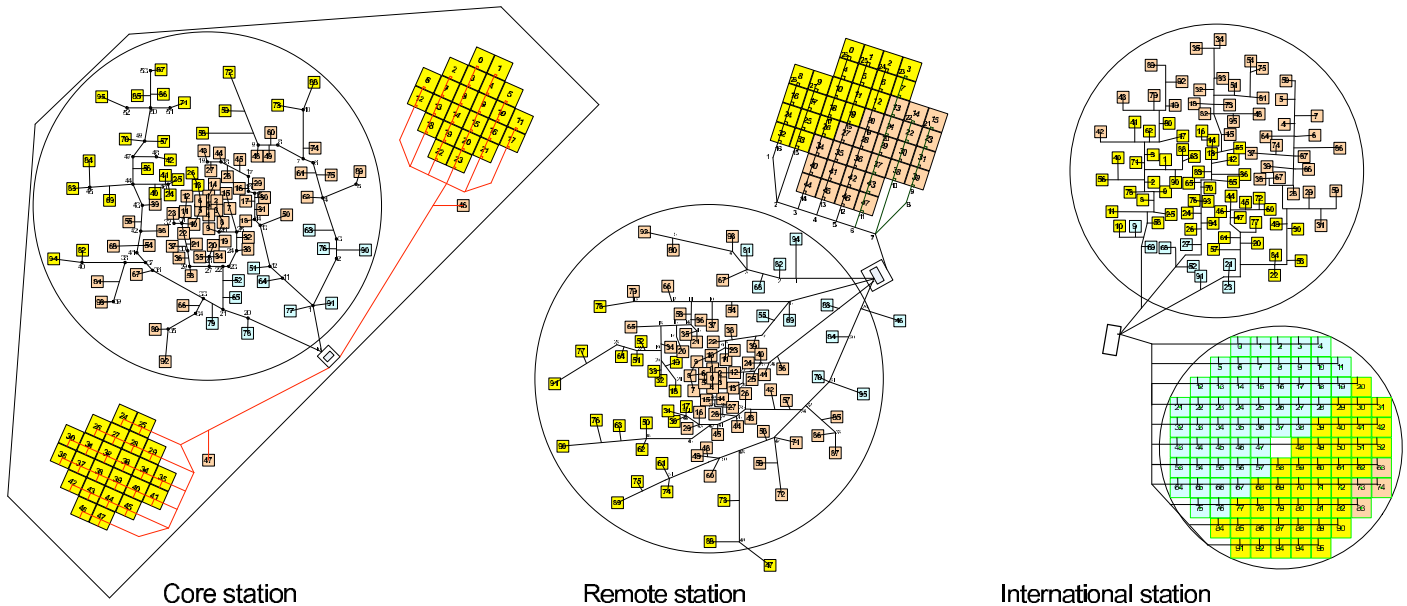


Figure 8: A schematic overview of the three different LOFAR stations (core, remote and international) from van Haarlem et al. (2013a). HBA antennas are denoted by the arrays of large square tiles, while the smaller squares represent dipoles in the LBA antennas. The stations layouts are not represented on the same spatial scale.

3.2 LOFAR antennas

If we want to measure any source in the sky, we need a tool to convert the incoming electromagnetic radiation in space to a quantifiable electric current. In the LOFAR telescope this is done by using a different combination of about $\sim 57\,000$ individual *dipole antennas*, or *dipole elements*. These typically consist of two symmetric linear wire elements, making it one of the simplest and cheapest ways to measure radio waves.

In LOFAR, the dipoles are categorized as either low-band antennas (LBA) or high-band antennas (HBA). LBA elements are dual-polarized single dipoles that operate between 30–80 MHz. Meanwhile, HBA elements consist of 16 dual-polarized dipoles combined in 4×4 square tiles that operate in the 110–250 MHz bandwidth range.

The LOFAR array utilizes LBAs and HBAs in three types of stations, depending on their distance to the center of the array: core stations (CS), remote stations (RS) and international stations. All of these stations have a different configuration of LBA and HBA elements, of which an overview can be seen in Figure 8. At the moment of writing, the full LOFAR array has a total of 24 core stations, 14 remote stations and 14 international stations.

3.2.1 Beamforming

The response, or sensitivity, of a dipole antenna to the sky is not constant. If placed horizontal to the ground, the dipole *radiation pattern*, or simply the *beam*, peaks at zenith and generally declines with $\sin^2(\theta)$ at angles closer to the horizon. The *half power beamwidth* (HPBW) describes the angular distance between the points at which the antenna’s sensitivity is 50% of the peak. Consequently, the HPBW determines the angular scale to which the antenna is sens-

itive.

By using analog and digital *beamforming* techniques, LOFAR can combine the signals of the individual LBA and HBA elements to form a *phased array*, with a combined beam response. Pointing the telescope is then done by adjusting (i.e. delaying) the signal phases in such a way as to create a beam in the target direction. Like with dishes, the HPBW is inversely proportional to the spatial separation between the antenna elements. Therefore, the HPBW of a combination of multiple antenna elements is always smaller than the HPBW of one of the individual antenna elements. To illustrate this, three examples of beamforming using LOFAR HBA elements can be seen in Figure 9.

3.3 Radio interferometry

As this thesis will discuss and frequently reference various interferometric techniques, it is important to provide a general introduction to the subject and highlight some of the most relevant topics one should know to understand this work.

Basics of interferometry

For the general explanative purposes in this section we assume the basic schematic of a two-element ($N_{\text{ant}} = 2$) interferometer as seen in Figure 10. This is relevant, as the synthesised response of a larger array ($N_{\text{ant}} > 2$) can roughly be described as the superposition of $N_{\text{ant}}(N_{\text{ant}} - 1)$ times the response of this setup. Most of the fundamental information in this section is based on the description by Thompson et al. (2017). For the discussion of specific properties we assume both antennas to be part of the LOFAR array as described by van Haarlem et al. (2013b).



Figure 9: Illustration of different beam responses that include elements of the HBA antenna in the LOFAR core. Relevant HBA elements are highlighted in white, while the corresponding beam can be seen in grey. The *left* image shows the beam of a single tile in the HBA array, the *middle* image shows three beams from the combination of all 24 HBA station tiles and the *right* image shows four simultaneous beams formed by combining HBA tiles from all 24 core stations (not all are displayed; [Stappers, B. W. et al. 2011](#)).

Each antenna in [Figure 10](#) is pointing at a source in the *far field*, i.e. the source is distant enough such that the arriving waves can be assumed to be planar. The spatial difference, defined as *baseline* length, between the two antennas causes a difference in the path length, and therefore arrival time, for the incoming waves. This time difference is called the *geometric delay* and can be expressed as:

$$\tau_g = \frac{d}{c} \sin \theta, \quad (3.4)$$

where d is the baseline length, c is the speed of light and θ is the angle between zenith and the pointing direction (also called the *phase center*). Additionally, the baseline length can be used to approximate the angular resolution of the array via:

$$\theta_{\text{res}} = \alpha \frac{\lambda}{d}, \quad (3.5)$$

where θ_{res} is the angular resolution, λ is the observing wavelength and α is a factor that can depend on a lot of different parameters (e.g. antenna configuration, beam shape, weighting scheme).

Response of an interferometer

The measured analogue wave signal by each antenna is filtered, amplified and digitized. Subsequently, a unique time offset (or *instrumental delay*) is added to compensate for the differences in geometric delay between each antenna. Finally, the two signals are combined by multiplication in the correlator and averaged over time. The final output response will be a sinusoidal fringe pattern of the form:

$$R \propto S_v \cos(\omega\tau_g), \quad (3.6)$$




Figure 10: Simplified diagram of a two-element interferometer pointing towards an incident radio wave (credit: [Paul Howland](#)).

where the amplitude is proportional to the flux density of the observed source (S_ν) and the phase depends on the observing frequency ($\omega = 2\pi\nu$) times the geometric delay. We can thus see that while the source moves over the sky (due to Earth's rotation), the geometric delay, and in turn the response of the interferometer, will continually change.

The FWHM of the fringes in the response is proportional to λ/d . This means that for the same observation wavelength, the fringes of large baselines will move quicker than those of shorter baselines across the same area of the sky. In turn this means that long baselines will be more sensitive to smaller scales in the observed source, while short baselines will be more sensitive to the larger scales.

For extended sources with a sky brightness distribution I_ν , we can integrate the response in Equation 3.6 over its sky area to obtain:

$$R \propto \int_{\text{source}} I_\nu(\hat{s}) \cos(2\pi\vec{d}_\lambda \cdot \hat{s}) d\Omega, \quad (3.7)$$

where \hat{s} is the unit vector pointing in the direction of the source and \vec{d}_λ is the baseline vector in units of λ .

The response in Equation 3.7 is the real (cosine) part of the signal. In a similar way we can obtain the imaginary part (sine) of the signal by using a second correlator that adds a phase delay of 90° . By combining these two signals using Euler's equation ($e^{i\phi} = \cos\phi + i\sin\phi$), we can define the response of our single-baseline interferometer to a source in the sky as a *complex visibility*:

$$\mathcal{V} \equiv V e^{-i\phi}, \quad \text{where} \quad V \propto S_\nu, \\ \phi \propto \vec{d}_\lambda \cdot \hat{s}.$$

Here V is the *visibility amplitude* and ϕ is the *visibility phase*. A visibility can be thought of as a measurement of the sky done by a single pair of antennas, thus giving information about the power on some projected angular size on the sky. The sampled angular size will depend on the fringe width, and thus the baseline length. For an extended source we can sum the individual visibilities over the source area to obtain:

$$\mathcal{V}_\nu = \int_{\text{source}} I_\nu(\hat{s}) e^{-2\pi i \vec{d}_\lambda \cdot \hat{s}} d\Omega. \quad (3.8)$$

Relationship between \mathcal{V}_ν and I_ν

To describe the response of our interferometer to any position in the sky, we can introduce a coordinate system for both the baseline vector, $\vec{d}_\lambda = (u, v, w)$, and the source, $\hat{s} = (l, m, \sqrt{1-l^2-m^2})$. The ground system has u pointing east, v pointing west and w pointing towards the source. In the sky system l and m are direction cosines of u and v , and thus point east and west, respectively. A drawing of the geometric relation between the two systems can be seen in Figure 11.




Figure 11: The relation between an interferometer array with baseline vector D_λ in the (u, v, w) plane and a source brightness distribution $I(l, m)$ in the (l, m) plane (Thompson et al., 2017).

Now, for illustrative purposes, we will assume that $w = 0$ for our interferometer array. This is theoretically only valid for small arrays and small fields of view/sources, but we will ignore this for the time being. With this assumption we can rewrite our interferometer response from Equation 3.8 to a 2-dimensional version:

$$\mathcal{V}_\nu(u, v) = \iint_{\text{source}} I_\nu(l, m) e^{-2\pi i (ul + vm)} dl dm. \quad (3.9)$$

This is a form of the van Cittert–Zernike theorem, which states that we can do a Fourier inversion of this result to obtain:

$$I_\nu(l, m) = \iint_{\text{source}} \mathcal{V}_\nu(u, v) e^{-2\pi i (ul + vm)} du dv. \quad (3.10)$$

Another way to write this is by using:

$$I_\nu(l, m) \longleftrightarrow \mathcal{V}_\nu(u, v), \quad (3.11)$$

where \longleftrightarrow denotes the Fourier transform. This is an important result, as we can now see that the visibility function $\mathcal{V}_\nu(u, v)$ can be related directly to the brightness distribution of the source $I_\nu(l, m)$ via a Fourier transform. Consequently, this means that to maximize our image fidelity,




Figure 12: The u - v coverage of the 8 hour long ILT observation of 3C330 used in this paper. Black represent visibilities with an international baseline, while orange signifies baselines shorter than $70\text{ k}\lambda$, i.e. baselines within the core and remote stations in the Dutch array.

we should make sure that our visibility function is completely sampled.

In reality we are not purely observing the brightness distribution and we should take into account two more issues. Firstly, our interferometer is not observing at a single frequency with an infinitesimal bandwidth, which means we have to include a bandwidth pattern $F_B(l, m)$ in our interferometer response. If we assume a uniform finite bandwidth of width $\Delta\nu$ (i.e. a top-hat filter), this will result in a $\text{sinc}(\Delta\nu\tau_g)$ envelope for our fringe pattern. Secondly, the brightness distribution of our source will get attenuated by the beam pattern $A(l, m)$ of the individual antennas in our array based on its distance from the phase center. These contributions are normally quite small relative to the contribution of the fringe pattern, and can thus in simple cases be neglected. We note that for LOFAR all stations are fixed, and thus also their response to the sky. This means that the main dependence of the final response is the observation angle from the horizon (i.e. declination).

Sampling of the Fourier plane

Each baseline vector in our interferometer array samples a discrete point in the Fourier, or, u - v plane. The values of u and v are also called the *spatial frequencies*, as they correspond to the frequency of the fringes (different from the observed frequency). The plane is always symmetric around the pointing center due to the existence of an equal and opposite baseline between every two antennas ($A \rightarrow B$,

$B \rightarrow A$).

The maximum number of visibilities or spatial frequencies measured in an interferometer array is limited by the number of individual baselines:

$$N_{\text{visibilities}} = N_{\text{baselines}} = \frac{N_{\text{ant}}(N_{\text{ant}} - 1)}{2}, \quad (3.12)$$

and we can plot them to obtain the u - v coverage of our telescope. The u - v coverage represents the non-zero area where our interferometer can make measurements. It is a good measure of the spatial structure that the telescope is sensitive to, and filling it in as complete as possible thus becomes essential to restoring an accurate source brightness distribution.

Large scale structure will be imaged by the smallest baselines, i.e. the innermost portion of the u - v coverage, while the small scale structure is probed by the largest baselines and thus outermost parts of the plane. Due to the nature of the baseline length distribution in most interferometric arrays, the large baselines are generally always under-represented with respect to the shorter baselines. Additionally, the restrictions of, for example, telescope arrangements and observation time will always leave us with an imperfectly sampled u - v plane.

The most common way to increase sampling in the u - v plane is by using *aperture synthesis*. Over the course of the observation the rotation of the Earth will cause the baseline vectors to rotate with respect to the source. This causes our visibilities to drift, causing elliptical tracks in the Fourier

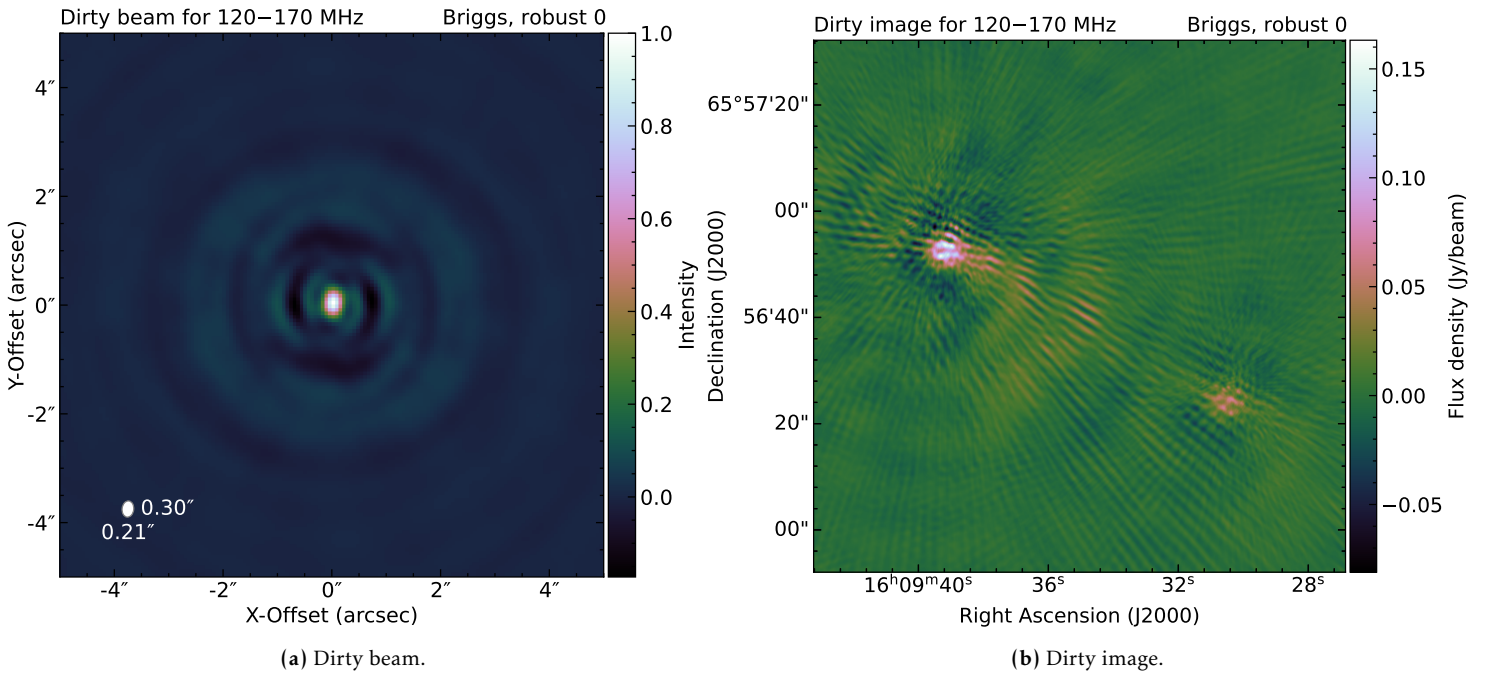


Figure 13: Cutouts of the dirty image and dirty beam of the partly calibrated ILT observation of 3C330. The dirty beam has major and minor axes of $0.30'' \times 0.21''$ and a position angle of -6.30° . The image has a noise value of $\sigma_{\text{rms}} = 3.13 \text{ mJy/beam}$ and a cell size of $0.0647''$.

plane. An example of such a coverage for our observation of 3C330 can be seen in Figure 12a.

Another way to fill in the intermittent space between tracks is to use *bandwidth synthesis*. By increasing the bandwidth frequency, the tracks will radially extend in a region corresponding to the largest and smallest frequency. We can see the added coverage that the 50 MHz bandwidth provides in Figure 12b.

3.3.1 Imaging and (de)convolution

Radio imaging constitutes the transformation of the observed visibility data to a 2D map of the corresponding sky brightness distribution. As described before, the core principle of imaging involves Fourier transforming our visibility function ($I_\nu \longleftrightarrow \mathcal{V}_\nu$). However, to complete this relation we have to take into account some other factors (Frey & Mosoni 2009, Thompson et al. 2017).

Firstly we have to take into account the fact that our telescope is not sensitive to all scales equally. From the filled in u - v plane we can define the sensitivity of the telescope to different spatial frequencies as the *spatial transfer function* $W(u, v)$. Secondly, we can apply a custom weighting to the visibilities represented by $w(u, v)$. A measured visibility then becomes:

$$\mathcal{V}_{\text{measured}}(u, v) = W(u, v)w(u, v)\mathcal{V}_\nu(u, v). \quad (3.13)$$

Using our relation in Equation 3.11 and the *convolution theorem* of Fourier transforms, we can see that our measured spectral brightness becomes:

$$I_{\text{measured}}(l, m) = I_\nu(l, m) ** B_0(l, m), \quad (3.14)$$

where $I_\nu(l, m) \longleftrightarrow \mathcal{V}_\nu(u, v)$, $B_0(l, m) \longleftrightarrow W(u, v)w(u, v)$ and $**$ denotes a two-dimensional convolution.

$B_0(l, m)$ is the synthesized beam of our interferometer, also called the *dirty beam*. It represents the response of the telescope to a point-source and can thus be compared to the PSF of optical telescopes. The missing measurements at certain interferometric spacings in our transfer function $W(u, v)$ will cause the dirty beam to have imaging artifacts. $I_{\text{measured}}(l, m)$ is the convolution of the dirty beam with the source brightness distribution and is therefore called the *dirty image*. Artifacts in the dirty beam will consequently also be represented in the dirty image. An example of the dirty beam for the 3C330 LOFAR observation can be seen in Figure 13a and the corresponding dirty image can be seen in Figure 13b.

Gaps in our u - v coverage can potentially be filled by an infinite number of visibilities, making it such that an infinite number of brightness distributions can actually fit to our visibility data. We can account and correct for the missing spacings in Fourier space by filling them in using constrained solutions based on all reasonable additional information. This process is called *deconvolution*.

The most widely used algorithm for deconvolution is the CLEAN algorithm (Högbom, 1974). Here the assumption is made that the source can be described by a finite number of model components, e.g. point sources (dirac-delta functions) or extended components of different scale sizes. Everything besides the source is assumed to be "empty" space.

The algorithm works iteratively by finding the brightest pixel in the image or, if given a mask, a user-defined region in the map. It fits a component to the found pixel and adds

this fitted component to an empty model. Subsequently the algorithm subtracts the model from the visibility data and then starts a new iteration. In this way a best-fit model is gradually built up. The process stops when either a certain iteration threshold is reached, or when the value of the peak has gone below a certain threshold value (typically set to around 3–5 times the rms noise value in the image).

A final *cleaned image* can be obtained by convolving the final model with a perfect Gaussian PSF and adding the result to the final residual map (which now ideally contains only the leftover noise).

At this stage the data is most likely still affected by various errors in phase and amplitude due to e.g. the effects described in Section 3.5. If the signal-to-noise ratio of the target data is bright enough, one can try and remove these errors by performing self-calibration on the target.

Another way to control the imaging process is by weighting the visibilities using the weighting function $w(u, v)$. The most widely used weighting schemes are called *natural* and *uniform* weightings. Additionally, the *Briggs* weighting scheme acts as a continuous combination between these two, where the *robust* parameter can be set to values from -2 (uniform) to 2 (natural).

Imaging of radio data is done by dividing the u - v plane into a grid of u - v cells. The amount of cells (i.e. pixels) and their angular size can be specified by the user and is typically chosen to be between 1/3 and 1/5 of the smallest beamwidth. The total weighting of each u - v cell is the sum of all the individual visibility weightings in that cell and the Fourier transform of the gridded plane is an image of the spectral brightness distribution.

Natural weighting is the most simple weighting scheme that weights each visibility data point by the inverse of its own rms noise variance ($1/\sigma_{\text{rms}}^2$). Cells with a high visibility density will thus naturally have a higher weighting. Consequently, imaging using this weighting scheme will generally result in images with the highest signal-to-noise ratio. Another consequence however, is that naturally weighted images will generally have the poorest resolution due to the radially declining density of visibilities in the u - v plane.

Uniform weighting is similar to natural, but with an additional step. Before Fourier transforming, all u - v cells are given a specific weighting factor such that the sum of all visibility weightings in each cell is equal to the cell with the highest summed weighting. This creates a uniformly weighted grid. Consequence of this is an increase in resolution, but also an increase of the rms noise in the image.

3.4 System performance and limitations

By being aware of the theoretical limitations of our telescope, we can be better prepared for potential upcoming issues during our final analysis. We will shortly discuss the most relevant limitations here.

3.4.1 Time and bandwidth smearing

Imaging assumes monochromatic data (a single frequency), but our observations have a finite bandwidth. The radial extension of the visibilities based on the frequency and the subsequent averaging will introduce *bandwidth smearing* in our image. This effect worsens when moving farther away from the phase center and can be minimized by making sure that:

$$\Delta\theta\Delta\nu \ll \theta_{\text{HPBW}}\nu, \quad (3.15)$$

where $\Delta\theta$ is the distance from the phase center, $\Delta\nu$ is the bandwidth, θ_{HPBW} is the beamwidth of the synthesized beam and ν is the observing frequency.

Bandwidth smearing causes sources to be stretched in the direction of the phase center and is especially apparent in wide-field images like those made with LOFAR. While this effect lowers the apparent surface brightness, the integrated flux density over the source will remain constant. To reduce the effect it is important to use small frequency chunks and move the phase center to our object of interest before we do any frequency averaging.

Time smearing happens when the time resolution of our data is not high enough, such that the source has moved about one synthesized beamwidth over the sky in between subsequent measurements. This effect is exacerbated by excessive averaging in time and can be minimized by making sure that:

$$\Delta\theta\Delta t \ll 1.37 \times 10^4 \text{ s} \times \theta_{\text{FWHM}}, \quad (3.16)$$

where Δt is the time resolution. Time smearing causes the sources to be stretched tangentially with respect to the phase center.

3.5 Calibration

The measured radio radiation observed with LOFAR is never a direct representation of the source. Both amplitude and phase of the signal have been affected by direction-dependent effects (DDEs) and direction-independent effects (DIEs). DDEs change depending on where the telescope is pointing, while DIEs are independent of telescope pointing position. The former is mainly governed by the propagation of the radio waves through intervening media, while the latter mostly consists of instrumental effects that have been introduced once measured here on Earth.

To ensure a high quality of our scientific results, it is very important to correct for these effects accordingly. In this section we will address the most important effects that we should take into consideration when observing with LOFAR and describe how we can correct for them. A more extensive overview of the effects of the most relevant neutral and ionized media on radio interferometry data can be read in [Thompson et al. \(2017, Sect. 13 & 14\)](#) and an excellent overview of the calibration of DDEs in specific can be seen in [Smirnov \(2011\)](#).

3.5.1 Direction-dependent effects (DDEs)

The most important DDEs on radio waves arise from the propagation through neutral media like Earth's troposphere, and ionized media like the ionosphere, the interplanetary medium and the interstellar medium. Of these, the Earth's ionosphere is by far the most troublesome for observations with LOFAR.

The ionosphere is the upper region of Earth's atmosphere that stretches from about 80–900 km above sea level. Gasses in the ionosphere are constantly being ionized by solar radiation, making it a medium with a high electron density that varies in both position and time. In contrast to the neutral medium, the ionized nature of the ionosphere makes it highly affected by the Earth's magnetic field. As a consequence of this, most of the ionospheric effects scale as ν^{-2} and are thus very dominant in the low frequency range.

Early low frequency (LF; ≤ 300 MHz) astronomy mostly ignored the impact of the ionosphere due to poor resolution, poor sensitivity and inadequate computing power to deal with the calibration. The high resolution and sensitivity of LOFAR-VLBI makes it so that ionospheric DDEs can no longer be ignored, and that these effects actually form the limiting factor in reaching the theoretical sensitivity. [Intema \(2009\)](#) provides a thorough overview of the ionosphere and its effects on LF signal propagation. Here we take a look at the most relevant effects for LOFAR.

Ionospheric refraction and phase delay: When a radio signal propagates through the ionosphere, it introduces a phase delay due to the added path length from refraction. When we for a moment neglect the magnetic field and consider a simple case of a linearly polarized monochromatic plane wave entering a cold plasma, the corresponding *refractive index* of such plasma can be described as:

$$n = \sqrt{1 - \frac{\nu_p^2}{\nu^2}}, \quad (3.17)$$

where ν is the radio frequency and ν_p is the *plasma frequency* given by

$$\nu_p = \frac{e}{2\pi} \sqrt{\frac{n_e}{\epsilon_0 m_e}} \approx 9\sqrt{n_e}, \quad (3.18)$$

where n_e is the electron density in m^{-3} , e is the electron charge, m_e is the electron mass and ϵ_0 is the vacuum permittivity. Average plasma frequencies in the ionosphere typically lie below 12 MHz. Incident radio waves with $\nu < \nu_p$ are fully reflected, while waves with $\nu > \nu_p$ are refracted via Snell's Law. By integrating along a path l that follows the line of sight (LoS) and assuming $\nu \gg \nu_p$, the added ionospheric phase delay can be described as ([Intema 2009; Morabito et al. 2022](#)):

$$\phi^{\text{ion}} = \frac{c r_e}{4\pi} \frac{\text{TEC}}{\nu}, \quad (3.19)$$

where r_e is the electron radius and TEC stands for the *total electron content* along the LoS in units of 10^{16}m^{-2} . By measuring the TEC value we can thus accurately describe the phase behaviour as a function of frequency and correct for it accordingly.

Faraday rotation: A second dominant factor on polarized radio emission passing through a ionized medium is a rotation of the wave's polarization plane due to the presence of a magnetic field. This phenomenon is called *Faraday rotation*. The angle by which the wave is rotated can generally be described as:

$$\Delta\psi = \frac{c^2}{\nu^2} RM, \quad (3.20)$$

where RM is the *rotation measure* ($RM \propto n_e B$). Due to the ν^{-2} behaviour Faraday rotation is dominant at low frequencies. By measuring the RM to the target one can correct for this rotation.

3.5.2 Instrumental effects (DIEs)

Most of the DIEs can be described as instrumental effects that are independent of the pointing direction.

Bandpass: Bandpass calibration entails an amplitude calibration per station with solutions that were obtained from observing a well-known and bright flux calibrator source. In this way station gains are corrected for both as a function of frequency and time and all stations will be on the same relative flux density scale. This is especially important for LOFAR, as the international stations have twice the amount of dipoles in comparison to the core and remote stations (96 instead of 48), leading to increase of factor 2–3 in sensitivity ([Morabito et al., 2022](#)).

Clock offsets and drifts: Each LOFAR station has a GPS-corrected rubidium *clock* that measures the time of an incoming signal. Remote and international station clock times are not synchronized with the single core clock used by the core stations, which can lead to offsets of about 100–250 ns ([Williams et al., 2016a](#)). These offsets cause significant phase delays, but seem constant over time and are independent of frequency. Therefore, it is generally sufficient to transfer a single median clock offset per station. Residual delays after correction can be attributed to clock drifts. These are time-variant offsets with typical values of ≤ 15 ns.

4 Observation and data processing

The observation used in this paper is part of the LOFAR Two-metre Sky Survey (LoTSS; Shimwell et al. 2017), which is an ongoing survey that aims to observe the entire northern sky in the 120-168 MHz range with LOFAR (van Haarlem et al., 2013b). This frequency range makes the LoTSS survey the highest frequency wide-area LOFAR survey project that is currently ongoing.

As this is an enormous project the observations in the LoTSS have been diligently planned to already span over a decade, having started in 2013 and continuing on in 2023 and beyond. Currently about 56% of the planned observations of the entire northern sky have been made. Meanwhile the already observed data has been made publicly available in chunks called data releases (DR). Besides the preliminary data release (Shimwell et al., 2017), there has been the first data release (Shimwell et al., 2019) and most recently the second data release (Shimwell et al., 2022), complemented with respective papers that describe a qualitative analysis about their general properties and characteristics.

For the purposes of the LoTSS survey, the northern sky is divided in a total of 3168 slightly overlapping circular fields or mosaics with a radius of 2 degrees (area of ~ 12.5 square degrees) that together cover the entire northern hemisphere. Of these 3168 planned LoTSS pointings, DR1 includes 63 (about 2 % of the northern sky) and DR2 includes a total of 841 (about 30% of the northern sky).

4.1 Observation details

The data used in this project was included in the second data release of the LoTSS survey (LoTSS-DR2). This release covers a total of 5634 square degrees in the northern sky and with a derived catalogue of about 4.4 million radio sources it is by far the largest release of the LoTSS project to date.

The observation of the data in this paper has been made using the same standard observation strategy that is used for the majority of the LoTSS sources (described in detail by Shimwell et al. 2017). The strategy includes an 8-hour on-source observation of an LoTSS field with the International LOFAR Telescope High Band Antenna (ILT HBA). The stations used for most of these observations are a total of 62 Dutch stations (2×24 core stations + 14 remote stations) and 13 international stations (DE601, DE602, DE603, DE604, DE605, FR606, SE607, UK608, DE609, PL610, PL611, PL612 and IE613), summing up to a total of 75 individual stations. A layout of the configuration can be seen in Figure 14. The mode that is run is HBA_DUAL_INNER, meaning that the dipoles in each core station are divided into two substations with a higher field of view (FWHM $\sim 4^\circ$) than the remote stations. Besides the source itself, two additional observations are made of standard flux density calibrators on either side of the field.

Each observation has a bandwidth of 48 MHz centered around the central frequency of 144 MHz (i.e. 120–168 MHz). This frequency range is chosen such that it avoids

Table 1: Observation parameters: ¹International LOFAR Telescope, ²Averaged to 12.205 kHz on the LTA, ³Used on the LTA website and part of Cycle 10 of the LoTSS observations, ⁴Given in J2000 right ascension (RA) and declination (DEC).

Observation properties		
Telescope ¹	ILT	
Observation date	15 June 2019	
Mode	HBA_DUAL_INNER	
Core stations	2×24	
Remote stations	14	
International stations	13	
Number of unique baselines	5550	
Time resolution	1 s	
Frequency resolution ²	3.05 kHz	
Central frequency	144 MHz	
Bandwidth	48 MHz	
Project name ³	LT10_010	
Target field		
Field name	P241+65	
Phase center ⁴	$16^{\text{h}}07^{\text{m}}45.7^{\text{s}}$ $+65^\circ 8'40.2''$	
Observation ID	L723262	
Time on source	8 hours	
Field calibrator		
Calibrator name	3C196	
Phase center ⁴	$08^{\text{h}}13^{\text{m}}36.0^{\text{s}}$ $+48^\circ 13'2.6''$	
Observation ID	L723282	
Time on calibrator	10 minutes	

the worst radio frequency interference (RFI) that exists near the higher part of the band and also the relatively high system equivalent flux density (SEFD) that occupies the lower end of the band. The data are taken with a time resolution of 1 second and a frequency resolution of 3.05 kHz per channel. This channel width is chosen such that AOFagger (Offringa et al., 2012) can more accurately exclude RFI. After this RFI excision, the data is averaged in frequency by a factor of 4 to 12.205 kHz per channel. The data is then uploaded to the LOFAR Long Term Archive (LTA) in the form of 244 subbands with 16 channels each.

4.2 Obtaining the data

The dataset including 3C330 lies in the mosaic called P241+65, named after its location in the sky (right ascension and declination). Using the Status Observation Tool¹, the corresponding observation ID for this field is retraced as L723262. From this tool we also obtain some additional information about the date of observation and how many international stations were observing. This tool also gives us the name and the ID of the main flux density calibrator that was used and when it was observed (in this case about 2 hours before the target was observed). The most relevant information about the observation can be found in Table 1.

¹<https://www.lofar-surveys.org/observations.html>

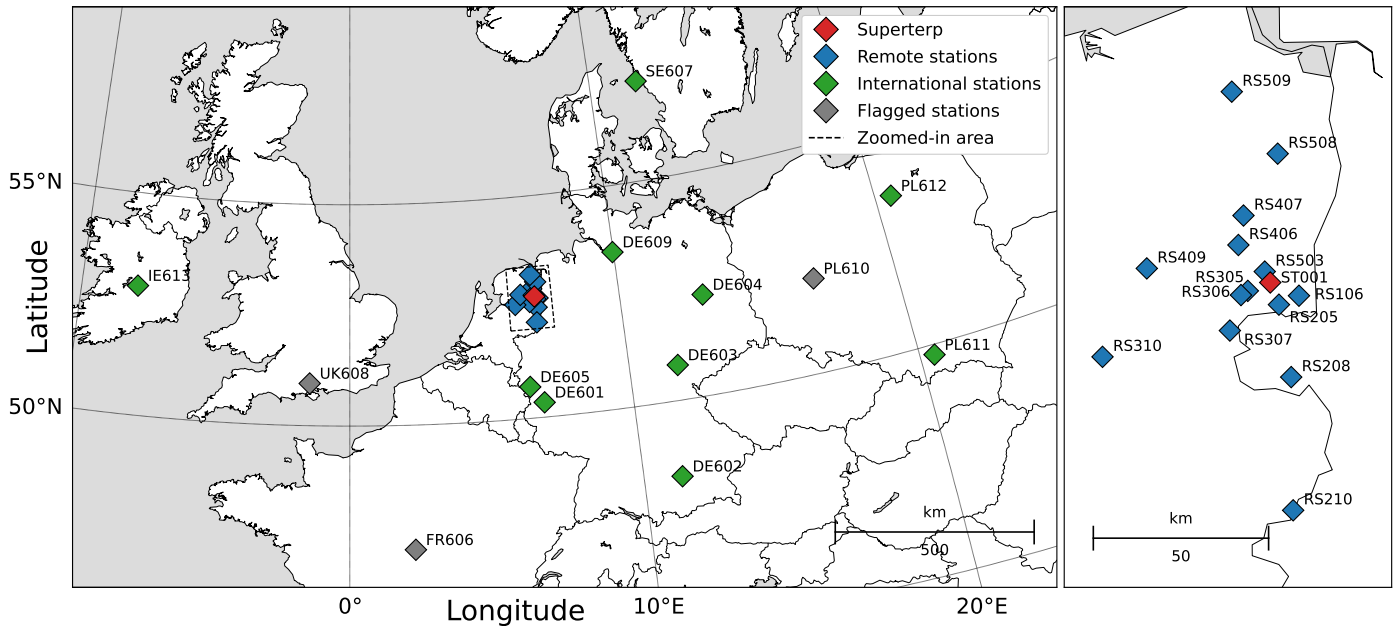


Figure 14: *Left:* a map of the total ILT configuration used for this observation. During the pipeline three international stations were fully flagged (UK608, FR606, PL610). *Right:* a zoom-in on the Dutch array. Core stations (CSs) are included in ST001.

The observation tool provides us with only one flux calibrator, but as noted before there are actually two flux density calibrators observed for each field. The reason for this is that the second calibrator can be used as a back-up in case something with the observation of the first calibrator might have gone wrong or in case the first is just not adequate.

Using the correct target and calibrator IDs, the observational data could be obtained from the LTA archive through the use of the LTA web service². The observation project which containing the data is named *The LOFAR Two-metre Sky Survey: Opening up a new window on the Universe* and falls under the LoTSS Cycle 10 projects. The corresponding project ID is LT010_10 and the data was released on 20 November, 2019.

When making a request for staging the data we make sure to only select data products (i.e. subbands) with a central frequency lower than 168 MHz. This is done partly because, as previously stated, subbands above 168 MHz are severely impacted by RFI, but also because these subbands do not have contiguous frequency bands and are mostly used for spectral line imaging. In the end this leaves us with a total of 231 data products for both the target field (P241+65) and the corresponding calibrator (3C196).

With the correct data selected, a data retrieval request was scheduled for staging and was assigned the ID number 68895. After some time our request was accepted and through the use of some uniquely provided .txt files (containing the appropriate commands) the data could be downloaded in form of an abundance of .tar files to our local system. A full, uncompressed dataset (target and calibrator) takes up about 16 TB of disk space, but when utilizing DYSCO compression (Offringa, 2016) this can be

lowered to about 4 TB. In our case the data was DYSCO compressed.

The initial downloading and processing of the data was done remotely on the Vaalserberg supercluster at ASTRON. Downloading and extraction of the data products took several hours and was done using Python scripts. In the end, the combined raw data of both the target and the source take up about 4 TB of disk space and are ready for further processing.

4.3 Data processing requirements

For the processing and calibration of our raw data we make use of the LOFAR long baseline pipeline³ (Morabito et al., 2022). This is a standardized calibration strategy that is specifically developed to calibrate the LOFAR HBA data from the sources in the LoTSS observed with the ILT.

Software requirements The long baseline pipeline makes use of specific versions of existing software. For this reason the recommended way to run the pipeline is by using the Singularity software (SingularityCE, 2021). Singularity makes it possible to create carefully configured working environments called images that can be shared with others, making it very easy to reproduce results of the same experiment. While the process of building an image is quite straightforward, for the purposes of this paper we have made use the image used by Kukreti et al. (2022), which has very similar goals to ours.

Running and interacting with parts of the pipeline is done using the genericpipeline framework⁴, which is part

²<https://lta.lofar.eu/>

³<https://github.com/lmorabit/lofar-vlbi>

⁴<https://www.astron.nl/citt/genericpipeline/>

of the LOFAR software package included in our Singularity image. In practice is to a Python script `genericpipeline.py` that is used to run `.parset` files containing specific user-defined steps and configurations. Such `.parset` files run non-interactively and, if for some reason interrupted, can be run again such that they re-start from where they left off.

Hardware requirements The recommended and tested hardware requirements for running the pipeline can be summarized as follows:

- 2x 16 core processor (32 threads) running at 2.10 GHz.
- 192 GB RAM (Random Access Memory).
- FDR Infiniband (i.e. a high throughput, low-latency connection).
- 100 TB disk space.

Unfortunately not all of these requirements could be met and the main differences between the available setup and the recommended one is that instead of two processing nodes we had access to one, we had about 25 TB of disk space available and instead of an Infiniband connection we had access to a 10 Gbps ethernet connection. While these limitations may have lead to increased computation times, in all cases the pipeline was still able to successfully run.

4.4 PREFACTOR pipeline

Preprocessing of the target and calibrator data is done using the PREFACTOR pipeline⁵ (as of recent renamed to the LOFAR Initial Calibration (LINC) pipeline⁶). The calibration strategy of the pipeline makes use of the Facet calibration scheme (van Weeren et al. 2016, Williams et al. 2016b) and is described in detail by de Gasperin et al. (2019). The main function of PREFACTOR is correcting either LOFAR HBA or LOFAR LBA datasets for various instrumental and ionospheric effects. The calibration is done in two successive `.parset` scripts, `Pre-Facet-Calibrator.parset` (for deriving corrections of DIES using the field calibrator 3C196) and `Pre-Facet-Target.parset` (for applying the derived DIE corrections and performing an initial self-calibration using a global sky model).

One of the tasks of the pipeline `Pre-Facet-Calibrator.parset` is to set the absolute flux density scale by finding the gain solutions from a model of the corresponding standard flux density calibrator. For this reason we run commit 79116d7 from the master branch, which includes the high resolution model of 3C196 that we require. On our setup the total runtime of this `parset` is 6 hours.

The first steps of the calibrator pipeline involve flagging and averaging of the data. This involves the advanced flagging of leftover RFI, known problematic or bad antennas and periods of time where the calibrator is below and elevation of 20° and averages the data to a time resolution of




Figure 15: Bandpass solutions in the XX polarisation as a function of amplitude against the frequency (120–168 MHz). The solutions that start above 300 represent the international stations, while the values of the Dutch stations are clustered between about 100 and 200.

4s and total of four channels per subband. After this, several solve steps are performed that address the following effects in order: correcting polarisation alignment between XX and YY, estimating the Faraday rotation measure, amplitude bandpass calibration, removing instrumental (clock) delays between core and remote station, calculating the ionospheric total electron content (TEC) delay values and phase calibration using the global sky model of 3C196. All direction-independent effects (polarisation alignment, bandpass calibration and clock delays) can be applied to the target field directly.

To check the quality of the solutions we can check the diagnostic plots that were generated by the pipeline (see Appendix A.1). In Figure 15 we see one of the diagnostic plots of the bandpass solutions for the XX polarisation. These amplitude values represent the factor by which the data has to be multiplied to get it in the units of Jy. From the plot it can be clearly seen that the international stations have values about 2-3 times higher than the Dutch stations. This can be explained by the facts that the international stations have more physical tiles and therefore more sensitivity and that the bandpass amplitude roughly reflects the antenna gain.

In Figure 16 we see the phase solutions in the XX polarisation for all stations in the observation. The uniformity in color of the core stations shows us that systematic effects were removed with high accuracy. We also see here that the phases of the remote and particularly the international stations vary a lot quicker with frequency and time than the core stations.

Applying the found solutions to the target field P241+65 is done by running the pipeline `Pre-Facet-Target.parset`. This will only apply the direction-independent solutions to the Dutch stations, and neglect

⁵<https://github.com/lofar-astron/prefactor>

⁶<https://git.astron.nl/RD/LINC>




Figure 16: A mosaic of the phase solutions in radians for all stations in the XX polarisation as a function of frequency (120–168 MHz) against time (0–10 minutes). Colours go from $-\pi$ (blue) to $+\pi$ (red). The first subplot represents CS001HBA, which is used as the reference antenna. The next 47 subplots represent the core stations and the last 27 subplots represent a combination of the remote and international stations. International stations are labeled by their country.

the international stations for now. One of the main reasons that we have to correct the Dutch and the international stations separately is due to the fact that there is a big difference in ionospheric effects for each one of these stations. Due to the wide geographical spread each one of the international stations sees through a very different part of the atmosphere, and as ionospheric effects become increasingly disruptive at lower frequencies it is necessary to calibrate them separate from the Dutch array. Additionally the international stations also have different clock drifts and beam sizes when compared to the core and remote stations.

The `Pre-Facet-Target.parset` pipeline first removes the international stations by flagging, which reduces the baseline amount and data size by about factor of 6. The next step involves demixing, which involves the removal of the brightest off-axis sources via phase-rotating, modelling and then subtraction of said sources. Corrections to the target field are then applied using the previously found solutions. In order, these are correction of polarisation alignment, amplitude bandpass calibration, clock delays, beam correction and correcting for rotation measure. In the end the field is phase-only self-calibrated against a downloaded sky model from the TGSS-ADR1 project (Intema et al., 2017), which is an all-sky radio survey at 150 MHz with a resolution of 25". An uninterrupted run of this pipeline on our system took about 36 hours.

Assessment of the calibration can again be done via the generated diagnostic plots (see Appendix A.2). In Figure 17 we see the phase solutions from self-calibration against the sky model. Here it can be seen that most core stations are quite stable in phase, but still experience some fluctuations mainly at the beginning of the observation. Another note is that most core stations do not center around a phase of 0, which at this plot scale would maybe be expected (phase fluctuations for the CSs are relatively small). For the remote stations we can see wide fluctuations across the entire phase range. Stations with the poorest solutions (noisy features) are RS210 and RS509 and partly poor solutions are found for RS310 and RS409.

Figure 18 can be used to identify stations with discrepant XX and YY polarization solutions. All core stations and most remote stations seem to have coherent solutions, while we notice the largest discrepancies for stations RS210, RS310, RS409 and RS508 and RS509.

Besides diagnostic plots the pipeline outputs a log file of badly behaved antennas that should be flagged for the next step. Despite the poor solutions for some antennas, no bad antennas were suggested for flagging in this file. For this reason we start by using all antennas in the forthcoming steps, justified by the fact that we can always go back to this step and flag said antennas if future problems arise.

4.5 LOFAR-VLBI pipeline

With the calibration of the Dutch stations completed, we now focus on the correction of the international stations. This correction is done with the LOFAR-

VLBI pipeline (Morabito et al., 2022). The `Delay-Calibration.parset` is run for this purpose.

This pipeline first applies the found direction-independent calibrator and target solutions from the `PREFACTOR` pipeline to the international stations. This creates a full dataset where all stations have been calibrated for clock, rotation measure, polarisation alignment and amplitude bandpass. Due to the `PREFACTOR` pipeline, the Dutch stations (CS and RS) have additionally been corrected for direction-independent phases. The data is then concatenated into chunks of 10 subbands, with a frequency width of $\Delta\nu = 1.95$ MHz each.

The next step involves crossmatching the LBCS and a LoTSS 6" catalogue (provided by Shimwell et al. 2022) to find the best in-field calibrator to solve for the direction-independent dispersive delays. This is done through an algorithm that aims to solve for TEC solutions with the least amount of scatter, favouring calibrators with high flux densities, compact structure and those that are located close to the phase centre. For the field P241+65 the pipeline found the best delay calibrator to be L422625 (3C330), which is also our target.

At this step in the pipeline process the datasets are extremely large (about 500 GB per 2 MHz chunk). As the next step involves self-calibration of the LBCS calibrator, a more appropriate and manageable dataset is required. This is achieved in a few successive steps:

1. The center of the field is phase-shifted to the coordinates of the delay calibrator 3C330 as noted in the LBCS catalogue.
2. The data is averaged by a factor of 8 in both time and frequency channels. This results in each subband going from 16 to 2 channels and having a total bandwidth of 390 kHz. Besides reducing datasize, this also reduced the field of view to limit contributions from nearby bright sources. The intensity losses at 30" from the phase centre are 0.1% and 2.3% due to time and frequency averaging respectively (using equations 18-43 and 18-24 of Bridle & Schwab 1999).
3. All core stations are combined into one super station called ST001. This makes all the interrelated baselines between the central stations (i.e. the bulk of the data) redundant.
4. All of the interrelated baselines are filtered out of the datasets.
5. All 2 MHz chunks are concatenated into one big list.
6. A beam correction or array factor is applied to account for the fact that all the individual beams of the central array vary in frequency and in time due to combining the tiles/dipoles together.

Together these steps result in a single measurement set that is around 57 GB in size, is centered on 3C330 and has a total of 260 subbands with each 2 channels to cover the bandwidth range from 121 to 168 MHz. The channel width is 97.656 kHz and the central frequency is 145 MHz.

With this dataset, 3C330 was subsequently used to solve for TEC (i.e. dispersive phase). Solutions were found using

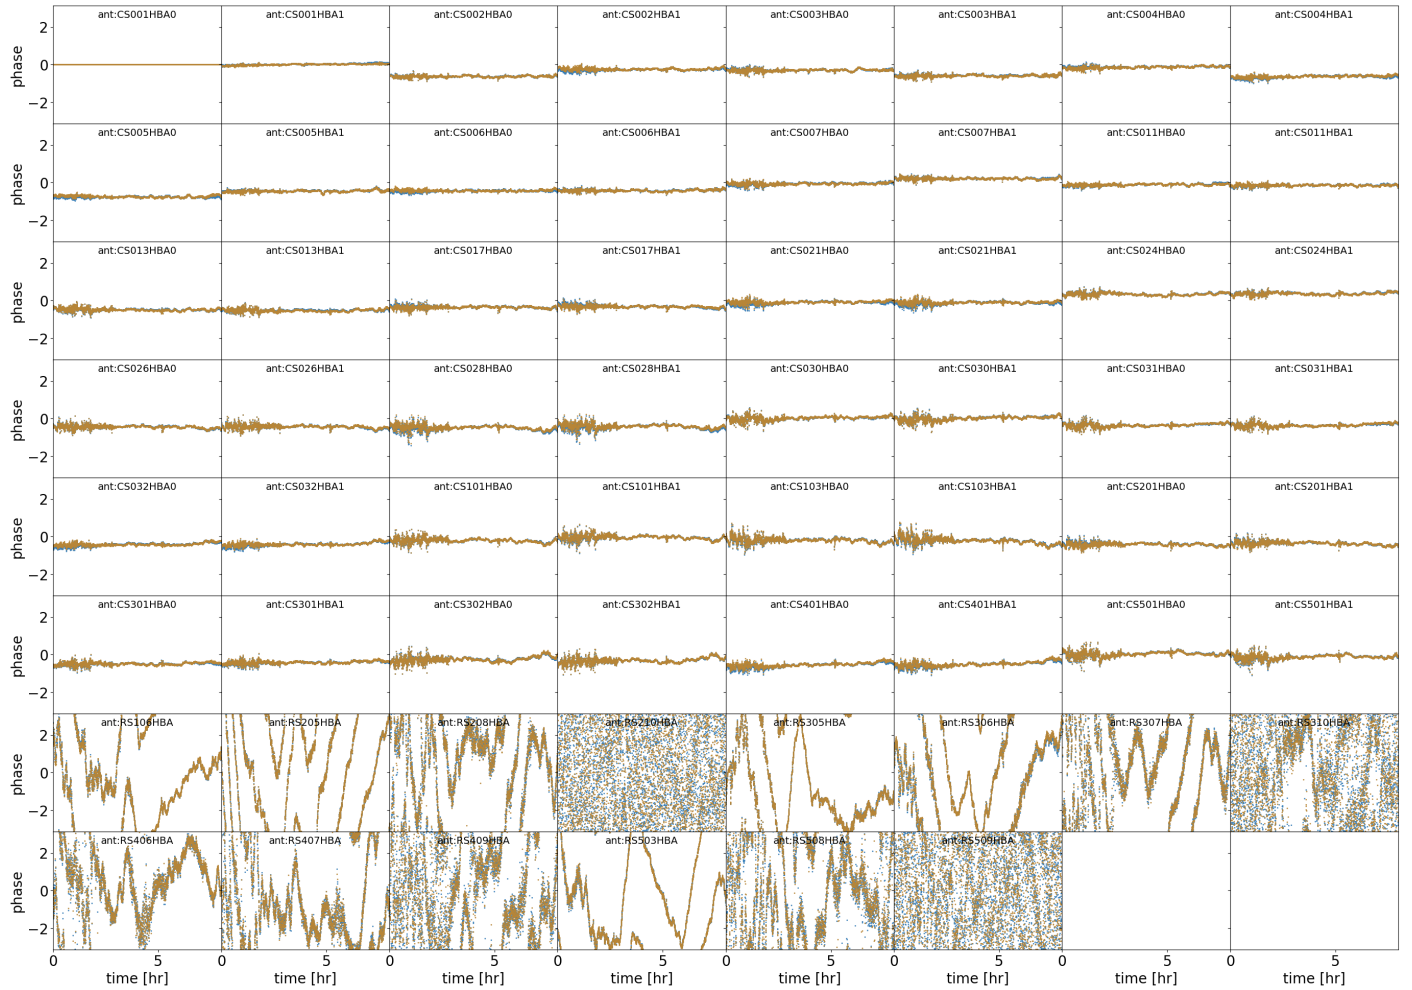


Figure 17: Mosaic of the phase solutions of the 2×24 CSs + the 18 RSs at 150 MHz obtained by self-calibration with the global sky model. Each plot shows the phase in radians (from $-\pi$ to $+\pi$) versus the observation time of 8 hours. The reference antenna CS001HBA can be seen in the top left corner and the bottom two rows represent the remote stations. Solutions for the XX and YY polarisations are plotted with different colors.




Figure 18: Mosaic of the difference in phase solutions between the two polarizations (XX – YY) of the Dutch stations at 138 MHz. Each plot shows the phase in radians (from $-\pi$ to $+\pi$) versus the observation time of 8 hours. The reference antenna CS001HBA can be seen in the top left corner and the bottom two rows represent the remote stations.

a point source model, a time resolution of $\Delta t = 16$ s and a frequency resolution of one subband ($\Delta\nu = 195$ kHz). Solutions are found by fitting curves to the phases over the full bandwidth (120–168 MHz). To avoid sudden jumps in the solutions, a phase offset is added to Equation 3.19 such that solutions are of the form:

$$\phi_{\text{ion}} = \phi_0 - \frac{cr_e}{4\pi} \frac{d\text{TEC}}{\nu}, \quad (4.1)$$

where ϕ_0 is the phase and dTEC is the differential TEC with respect to the reference antenna ST001. A single dTEC value describes the phase corrections at all frequencies.

To further improve the quality of the solutions for the international stations, the pipeline uses a $u-v$ cut of 50 k λ (~ 100 km) for the solving step. This results in an exclusion of the short baselines and therefore the sensitivity to the large scale structure and bright sources away from the pointing centre.

Inspections of the quality of the dTEC solutions is done by analyzing pipeline-generated diagnostic plots. Figure 19 shows an example of these solutions. It can be seen that dTEC values range between about -2.5 to 2.5 TEC for most stations. Average solutions for differential TEC of the remote stations ranges close to zero, which is expected as they have already been calibrated by PREFACTOR. Most notable exceptions are RS210 and RS509, which show quite some wide ranging noisy solutions. Additionally, some stations like RS307, RS310 and RS409 show some incoherent solutions in the first ~ 3 hours of the observation.

Taking a look at some of the international stations, we can see that there is a dTEC span of about -5 to 5 TEC. Solutions for the two Polish stations PL611 and PL612 go to -5 in the middle of the observation, where they seem to bottom out for a period of about 2 hours. The opposite happens for the Irish station IE613, which seems to top out at a dTEC of 5 for a period of about 3 hours. Moreover, all of these solutions seem to not be very smooth and exhibit a certain degree of noise. The Swedish station SE607 seems to have some incoherent solutions in the first three hours, but smooths out thereafter. We note that stations FR606, UK608 and PL610 were automatically flagged by the pipeline, explaining their flat solutions.

Another way the pipeline checks the quality of the TEC solutions is by applying them to the dataset and subsequently imaging the source. This is done in the pipeline using DIFMAP. Additionally, a self-calibration routine is performed to try and further improve image fidelity. As the script expects resolved calibrators, the data is phase-shifted to the brightest pixel. Figure 20 shows the generated pipeline image at full resolution. An additional cleaned map with a $u-v$ cut of 50 k λ is generated (see Figure 37 in Appendix A.3), which essentially excludes the international baselines and only uses the Dutch array. This allows for a quantitative comparison to the calibrated LoTSS data (see Figure 6).

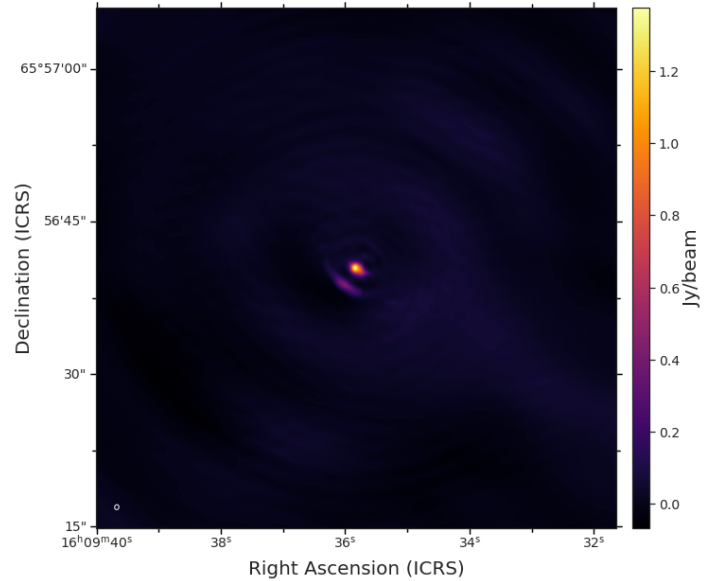


Figure 20: Automatically generated pipeline image of 3C330 with the full resolution (i.e. using all available baselines). The phase centre is positioned at the location of the N-W lobe. The S-E lobe cannot be seen due to the mapsize.

In this stage a number of methods were tested to try and improve the TEC solutions. These methods include flagging individual stations (e.g. RS210, PL610), changing calibration algorithms (DDEcal instead of gaincal) and changing the solution intervals (4s, 8s, 32s). Unfortunately, none of these methods showed significant improvements and in most cases, actually decreased solution coherence. The suspected cause for this is the point-source model, which was left unchanged until now. To create an accurate model we decide to end our work in the pipeline and continue locally. As such, the final steps of our data preparation includes applying the best found TEC solutions to our averaged dataset and converting it to a non-DYSC0- compressed measurement set.

4.6 8.5 GHz VLA Observation

Besides our ILT observation, we also used publicly available VLA data for the creation of our spectral index map. Using the NRAO Data Archive⁷, we have downloaded archival VLA X-band (8.5 GHz) data taken using the A-configuration (maximum antenna separation of 36 km). The dataset contains two alternating scans of both 3C330 and its respective phase calibrator, 1642+690, in two averaged down spectral windows. Additionally, scans of three different flux density calibrators (3C48, 3C137 and 0552+398) are provided. More observational details can be seen in Table 5 in Appendix B.2. This specific dataset was chosen for its similarity in angular resolution to our LOFAR-VLBI observation and because of its smaller size, allowing for quicker processing time.

⁷<https://data.nrao.edu/portals/>




Figure 19: Mosaic of the dispersive delay solutions for 3C330. Each plot shows the differential TEC (dTEC; from about -5 to 5) versus the observation time of 8 hours. The plot shows solutions for the 14 remote stations (top two rows and first two plots on the third row) followed by solutions of the 13 international stations. The reference antenna is ST001 (final plot).

5 Results

In this section we will show and describe what was done after the initial pipeline calibration. This entails manual flagging of the data, a further (self-)calibration of the target to create a multi-frequency model and creating flux density and spectral index maps of 3C330.

5.1 Spectral energy distribution (SED)

To obtain a general understanding of the flux density behaviour of 3C330, we will plot its spectral energy distribution by using flux density measurements from the literature. Additionally, we want to fit a model through these measurement points such that we can use it to flux calibrate our data in the $\nu = 121 - 166$ MHz frequency range.

To keep the measurements consistent we want to try and make sure that all flux densities are calibrated to the same flux scale. For this reason we will use the scale from Roger et al. (1973, RBC) at frequencies $\nu < 325$ MHz, as it is thought to be most consistent in this regime (Scaife & Heald (2012), Rees (1990a)). For frequencies $\nu > 325$ MHz we will use the scale from (Kellermann et al., 1969, KPW), as in this regime it is in agreement with the RBC scale.

A catalog of flux measurements was found using the photometry overview of 3C330 on the NASA/IPAC Extragalactic Database (NED⁸). We filtered out duplicate measurements and peak measurements and, if necessary, scaled the flux density to either the RBC or KPW scale using the appropriate factor. In the end we compiled a list of 21 unique observations at frequencies between 10 MHz and 14.9 GHz. An extensive overview of all of the used measurements, their references and their scaling factors can be seen in Table 4 in Appendix B.1.

The fitting of our measurement is done using the Bayesian fitting method model of Scaife & Heald (2012):

$$S [\text{Jy}] = A_0 \prod_{i=1}^N 10^{A_i \log^i[\nu/150 \text{ MHz}]}, \quad (5.1)$$

where S is the flux density in Jy, N is the polynomial order, A_i are the i -th order polynomial coefficients and ν is the frequency in MHz. This model is valid for frequencies between 30 and 300 MHz and was developed based on empirical SED data of six bright 3C sources (not including 3C330). More specifically, the main goal of this model's development was to function as an initial reference point model for future observations of similar objects with LOFAR.

The SED of 3C330 together with the best fit and residuals can be seen in Figure 21. For the fitting process we used literature values from 10 MHz to 14.9 GHz. The inclusion of data points outside of the consistent range (30–300 MHz) is necessary to constrict the edges and is also consistent with other fits by e.g. Scaife & Heald (2012) themselves, and McKean et al. (2016). We used an orthogonal distance regression method in Python (`scipy.odr`)

and found a second-order polynomial ($N = 2$) to best fit our data with values for the polynomial coefficients of:

$$\begin{aligned} A_0 &= 35.027 \pm 0.330, & A_2 &= -0.136 \pm 0.016, \\ A_1 &= -0.579 \pm 0.026. \end{aligned}$$

The fit has a residual variance of $\chi_{\text{red}}^2 = 4.44$, with most residuals coming from the low frequency regime ($\nu < 100$ MHz). Errors in the LOFAR HBA frequency regime range between 0.41% and 1.22% at 121 and 166 MHz respectively.

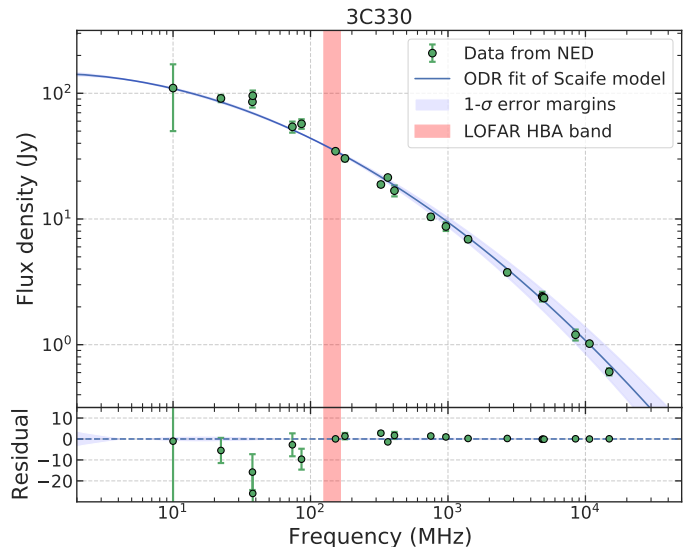


Figure 21: A log-log plot of the spectral energy distribution of 3C330. Literature values of the flux density as a function of frequency are fitted by a 2nd order polynomial model. Residual values (data - model) can be seen below.

To get an estimate of the overall spectral index of the source we will assume the region of $\nu = 121 - 166$ MHz to be linear. Using a positive sign convention ($S_\nu \propto \nu^\alpha$) we then calculate a straight-line spectral index of $\alpha_{166}^{121} = -0.573 \pm 0.026$.

5.2 Data inspection and flagging

Improving the calibration of the LOFAR-VLBI data requires a more accurate model than the LOFAR-VLBI pipeline can provide us. To analyze the cause for these limitations, and to get a better feel for the data, we therefore choose to continue our work locally with the averaged measurement set of 260 subbands.

By using the `plotms` task in CASA we can generate some diagnostic plots of our currently calibrated data, of which Figure 22 is an example (see Appendix A.4 for additional plots). We can see that our general visibility amplitudes behave as expected. The very short baselines detect a large amount of sources in the field. After about 200 k λ the baselines of the international stations start to slowly resolve these sources (including 3C330), which explains the flattening amplitude behaviour at the higher baselines.

⁸<http://ned.ipac.caltech.edu/>

High narrow peaks in the visibilities, indicate irregular data that needs to be flagged.

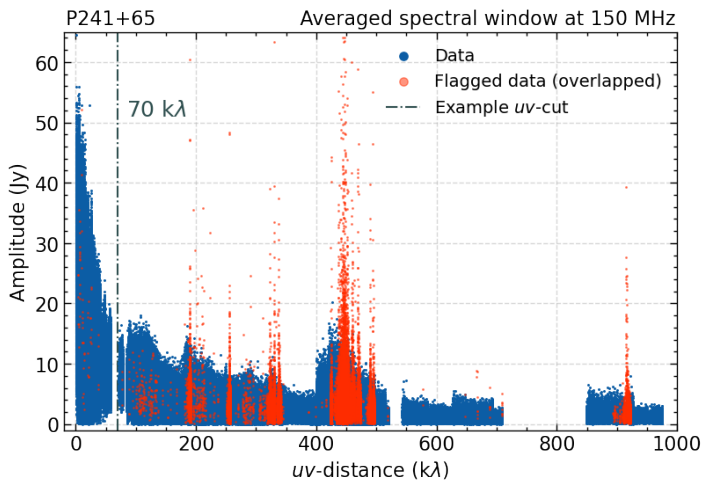


Figure 22: The visibility amplitudes versus the uv -distance (i.e. baseline length in number of wavelengths) of an averaged spectral window at 150 MHz. Blue data represents unflagged visibilities, while overlapped red data are the manually flagged visibilities. An example of a uv -cut at $70k\lambda$ can be seen.

Through thorough inspection we compile a list of all anomalous data present in our current dataset. This is done by very carefully isolating the origin (i.e. specific antennae, time ranges, baselines or channels) of these data, flagging said specific part and plotting again. This cumulative manual flagging strategy ends when the diagnostic plots show no more obvious bad data. We conclude by running a final script that contains the accumulation of all identified flags.

The most notable bad data comes from the station PL612, which shows numerous large visibility spikes starting about four hours into the observation and continuing for about two hours thereafter.

All German stations (DE601, DE602, DE603, DE604, DE605, DE609) have one specific time bin with extreme amplitudes. This bin is identified as an edge-channel of a larger time range that AOFFLAGGER has already flagged.

Some remote stations (RS106, RS406) show spiky time bins early in the observation and some (RS106, RS205) show some bad channels.

Besides station-wide flags, we identify a number of specific baselines that show bad features as either a function of time (noisy) or as a function of frequency (e.g. Gaussian-like amplitudes).

5.3 Self-calibrating individual subbands

After manual flagging we go on to further calibrate our data. Imaging of the general pipeline calibrated data up until now (see Figure 20) has revealed a clear detection of 3C330, indicating that the data is a good candidate for further self-calibration. Considering the size of the dataset and our locally available computing resources, we will use

a strategy of individually self-calibrating a subset of subbands spread out evenly across the frequency band. Working on one subband at a time reduces the size of the measurement set from 57 GB to 200 MB per measurement set and makes computing times a lot more manageable.

Our manual selection, or subset, includes subbands centered at the frequencies 121, 130, 138, 146, 150, 158 and 166 MHz. These bands are spaced evenly around 150 MHz ($\Delta\nu \approx 8$ GHz) and were manually inspected to make sure they are not dominated by flagged data or noise. Each subband has a total of 2 channels, a total bandwidth of 195.312 kHz and a time resolution of 8 s. To further decrease the file size we make sure to split out only the XX and YY correlations.

Our calibration strategy involves making an as accurate as possible model of one of the central frequency subbands and using this model for initial calibration of the adjacent subband(s) in a branch out fashion. This method allows for good efficiency, as we do not have to self-calibrate each subband from scratch. Additionally, it ensures that we iteratively build upon an overarching model instead of creating separate models each with their own independent calibration solutions. This is especially important as we are self-calibrating by hand, which increases the likelihood of divergent solutions when calibrating the subbands independently from one another. We should note however that this goes both ways, meaning that with our strategy a calibration error in the initial model could be transferred to an adjacent subband and ultimately become ingrained in the final model. To make sure this does not happen, extra care should thus be placed on the self-calibration of the first model.

Self-calibration of the subbands was done in DIFMAP (Shepherd, 1997, version 2.5e). The subband of 150 MHz was chosen as the starting model because of its central location in the bandwidth, although we note that the 146 MHz subband would also have sufficed for the same reason. A map size of 4096×4096 with a cell size of $0.05''$ ensures the entire source is captured with a good sampling (around 5 pixels in the beam). The corresponding dirty image can be seen in Figure 23a.

Applying a u - v cut Like the strategy of Morabito et al. (2022), we apply a u - v cut to the data before we start cleaning and calibrating. Because of the wide field-of-view of the shorter baselines, there are a large amount of sources outside of our map that we cannot account for (i.e. they are not represented in our model). We find that limiting our u - v distance to $> 70k\lambda$ cuts out all the interrelated baselines of the Dutch array, which removes the bulk of these sources. Additionally, it falls right in a gap in the u - v domain, minimizing the cutting out of unnecessary data (see Figure 22). It is important to note that while we create a model and generate calibration solutions with this u - v cut applied, we make sure to remove the limit when we apply the solutions to our visibility data. This ensures that all of the visibilities of all the stations are corrected based on an as accurate as possible model.

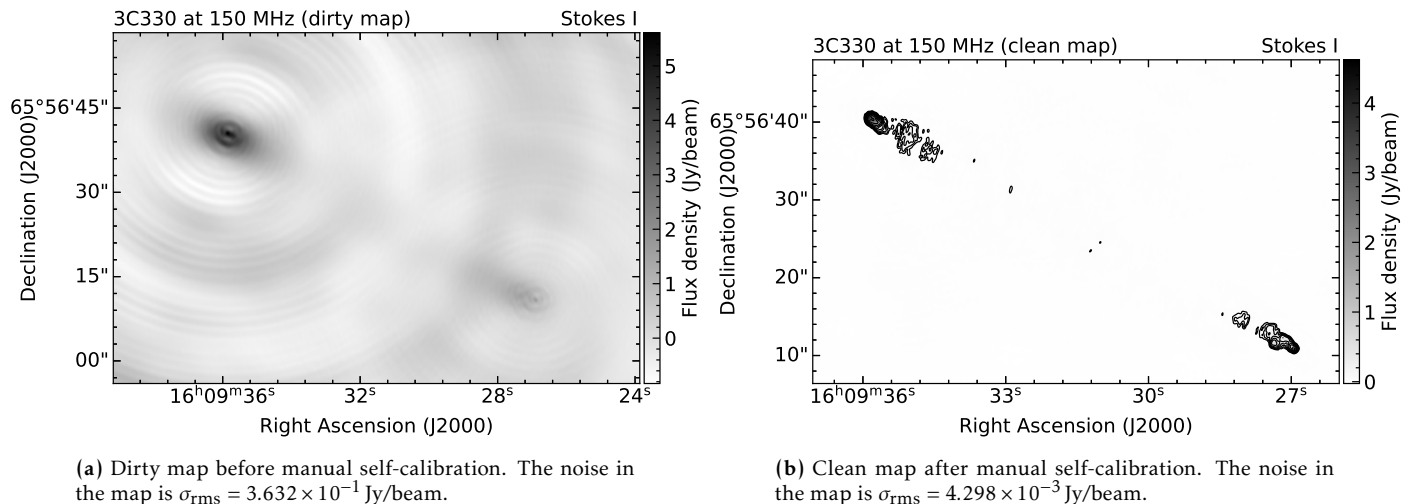


Figure 23: Maps of before and after manual self-calibration of one subband at 150 MHz. Both maps have beam sizes of $0.454'' \times 0.357''$ with a position angle of -6.95° . The combined core station ST001 has been flagged.

We tentatively start cleaning and phase calibrating with a point-source model on the NE hotspot. This is a slow process in which we make sure to only include flux that corresponds to the known structure of the source from e.g. Gilbert et al. (2004). We find best results by starting our phase calibration with a solution interval of $\Delta t = 64$ s and gradually decrease this to phase solutions of $\Delta t = 8$ s as S/N increases. We track improvements of the solutions by an increase in the ratio of peak flux density and the map noise. When improvements with phase calibration stagnate, we try an overall amplitude calibration (solution interval of $\Delta t = 900$ minutes) and afterwards amplitude calibrations with slowly decreasing solution intervals down to $\Delta t = 3$ minutes. This shortest interval seems sufficient to track the smallest changes in baseline amplitudes. The overall strategy is similar to the one described by Morabito et al. (2022, section 4.4). Amplitude correction is however done conservatively, as here bad calibration solutions could get baked in the data. An image of our final corrected data can be seen in 23b.

Our final model at 150 MHz was used for initial calibration of the subbands at 146 and 158 MHz. As the model is quite precise, we found a phase solution interval of $\Delta t = 8$ s to give the best results. After this initial calibration a deep clean was done to scrub away the excess or shortage of flux that is a natural result of the difference in total flux present at different frequencies (see Figure 21). Flux outside of the 150 MHz clean boxes was only added if it was a believable part of the source. After all residual flux in the mask was cleaned, another round of phase self-calibration and cleaning was done, repeating until no further improvements could be seen. To conclude the initial calibration, we perform an overall amplitude self-calibration and afterwards an amplitude calibration with a solution interval of 30 minutes. After this initial round of self-calibration of the data a new round of self-calibration was performed. The established model was deleted and more rounds of cleaning and self-calibrating were performed. Here the

choice of strategy and solution intervals remain the same to what was used before. In most cases this second clean/self-calibration iteration proved to result in a sufficiently accurate model.

Each newly obtained model was used as an initial model for the next adjacent subband. This entails using the 146 MHz model to calibrate the 138 MHz, using the 138 MHz model to calibrate the 130 MHz subband, and so forth. In this way each subband gets the most accurate initial model, i.e. the one that lies closest in terms of frequency. Following this method all seven subbands were calibrated.

5.3.1 Flux calibration

To correct for the relative changes in flux density that were introduced during self-calibration, we have to flux calibrate each subband separately against our fitted SED model (see Figure 5.1).

To get the total source flux $S_{\nu, \text{source}}$, we clean each subband in WSCLEAN. A mask of the source area is supplied to restrain the inclusion of sidelobe structure. Cleaning is done up until an automatic threshold of $3 \times \sigma_{\text{rms}}$. This threshold proved to be a sufficiently deep clean for all subbands. The source flux is then read-out from total flux in the model.

Calibrating our data is done using CASA. First we calculate the correction factor as:

$$x = \sqrt{\frac{S_{\nu, \text{model}}}{S_{\nu, \text{SED}}}}, \quad (5.2)$$

where $S_{\nu, \text{SED}}$ is the flux at the frequency ν taken from interpolating the fitted model in Figure 21. The square root is necessary as CASA squares the parameter due to XX and YY. An overview of the total fluxes and their correction factors can be seen in Table 2. To finalize, we run gencal to generate the appropriate calibration table for each correction factor, and subsequently run applycal to scale the flux in each subband by its respective correction factor. To

check if flux scaling was done correctly, we clean again with `WSCLEAN` using the same script.

Table 2: For each of the model frequencies the fluxes from both the model and the SED with their corresponding correction factor.

ν (MHz)	S_{model} (Jy)	S_{SED} (Jy)	x
121	25.52	39.60	0.803
130	27.41	38.00	0.849
138	28.51	36.70	0.881
146	35.06	35.58	0.993
150	32.05	34.89	0.958
158	31.27	33.98	0.959
166	31.25	33.01	0.973

5.4 Creating a concatenated (MT-)MFS model

The next step was multi-frequency synthesis (MFS) modeling, meaning combining or concatenating the seven subbands again and creating a new model based on the entire frequency range. For this step we have changed our working software from `DIFMAP` to `CASA`, as it is better equipped to handle the concatenation of data files. The calibrated data in `DIFMAP` are exported to UV FITS files (.uvf) and imported in `CASA` as MeasurementSet files (.ms). The flux calibrated subbands are subsequently concatenated with the task `concat`.

Similarly to the individual subband models, we create our MFS model with `tclean`. Unlike before, we now set the deconvolver to the mode `mtmfs` (multi-term multi-frequency synthesis). Unlike the standard `mfs` algorithm, `mtmfs` will attempt to fit a polynomial through the source spectrum. This is needed, as we are now covering a much wider bandwidth. The parameter `nterms` governs the polynomial order, and we set it to be `nterms` ≥ 2 (i.e. non-linear). Several runs with different orders were attempted, but no major differences were found). Another main parameter of choice is `scales`, which governs the sizes of the Gaussian components that the algorithm will try to fit to the data (in units of pixels). All visible (unresolved) features in the map have a shape and size very comparable to that of the synthesised beam. This means that with for example a synthesised beamsize of $0.35''$ by $0.25''$ and a cellsize of $0.0647''$, we can expect most of the features in the image to have a pixel size in the order of around 4 pixels.

Cleaning is done initially with uniform weighting and a $u-v$ -cut of > 80 k λ . For scale sizes we follow the `CASA` documentation and use scales that are 1x, 3x and 5x the beamsize. In addition to this, we use a scale size of 0 such that delta functions can also be used for cleaning. After no more source flux can be seen with uniform, we change to briggs and natural weighting to clean up any leftover larger flux components. Additionally, we slowly lower the $u-v$ -cut to include more of the central diffuse flux inbetween the bright lobes. We end by performing a round of phase self-calibration with a solution interval of 16s. Our final cleaned image of the seven subbands can be seen in Fig-

ure 24. The corresponding model and residuals can be seen in Figure 25 and Figure 26, respectively.

5.4.1 Calibrating all subbands

We can further our analysis by using our now accurate MT-MFS model to calibrate the full (averaged) dataset of 260 subbands. For this we only use the XX polarization for reasons stated earlier and to also increase the computing speed. We also make sure to apply the flag script that was made during our initial data inspection. As an additional measure we actually remove the phased-up array ST001 in this step (more on this in Section 6.3). The calibration of the data was done in `CASA`.

To try and optimize our results, several calibration strategies were tested. This entails a combination of several solution intervals (int/8s, 16s, 128s, 30min, inf) and calibration modes (p, a, ap) in the `gaincal` task. The resulting dirty images were quantitatively compared by their dynamic range (maximum/peak flux divided by the rms). From these comparisons, the best results were obtained from an amplitude and phase calibration (ap) with a solution interval of 8 seconds. Using this strategy, a total of 3595 solutions were found (100% of solution intervals) and individual antenna solutions below a S/N of 3 were flagged. These flags were also applied to the data in the `applycal` step thereafter. Additional settings for `gaincal` are an UV-cut of 50 k λ and RS407HBA as a reference antenna due to its good uptime, stable data and central position in the array.

After the initial calibration the data visibilities were inspected once more. Some of the edge subbands did not lie within the range (121–166 MHz) of the MT-MFS model and thus were not calibrated. Because of the edge subbands' naturally lower quality, and because this applied to only a small number of them (~ 10 subbands), they were manually discarded from the dataset. Besides this no significant new outliers or calibration errors were detected.

Due to the increase of data size, we clean our visibilities non-interactively using the more efficient `WSCLEAN`. For this we use the setting of multi-scale cleaning with Gaussian components. To limit the cleaning of any noise features we create and use a custom-made mask based on a $25 \times \sigma_{\text{rms}}$ clip of the 7 subband MT-MFS image. Additionally, we make sure to increase the auto-mask parameter to 10, although this might not be necessary. To limit the software's tendency to fit scale features that extend beyond the mask, we use a higher-than-default value of 0.7 for the `multi-scale-bias` parameter. The result can be seen in Figure 27. Integrated flux measurements of the individual components can be seen in Table 3.

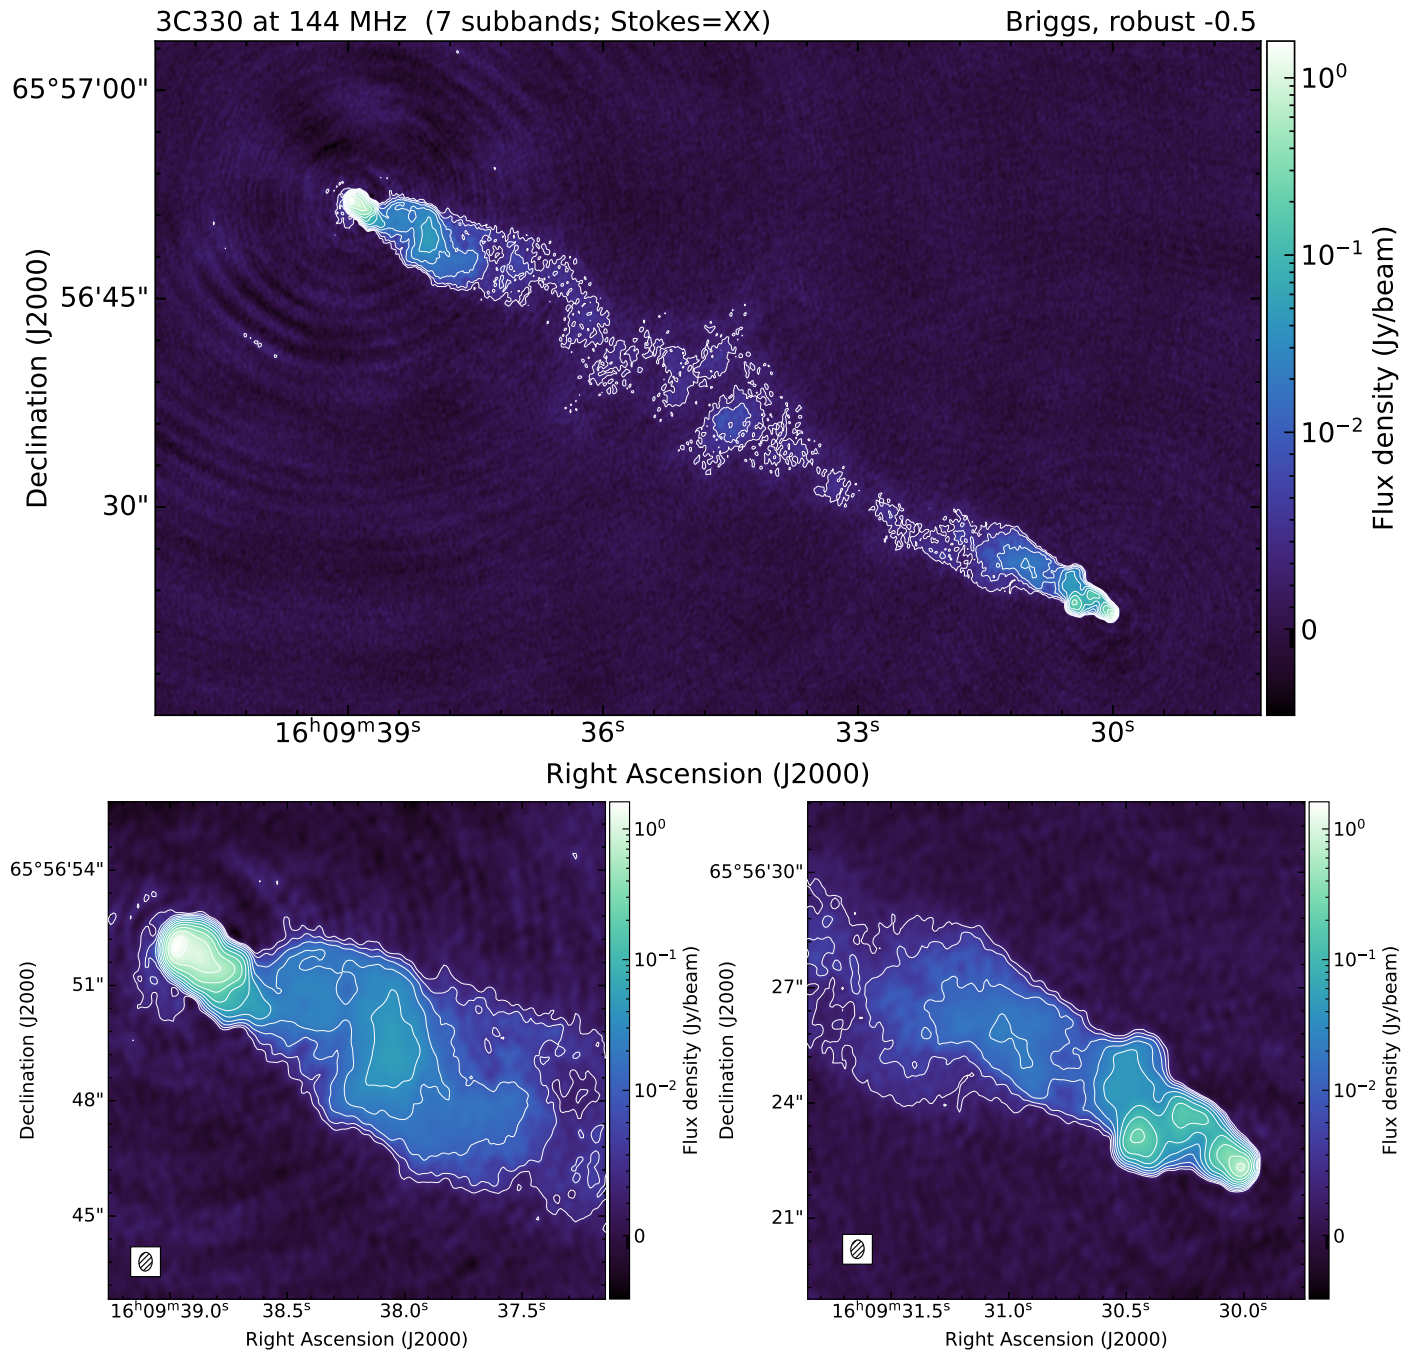


Figure 24: *Above:* flux density map of the full source 3C330 at 144 MHz using 7 subbands spanning 121–166 MHz. The map is convolved with a synthesised beam of $0.26'' \times 0.37''$ and a position angle of -4.77° . Contours are overlaid as $\sigma_{\text{rms}, 144\text{MHz}} (= 5.804 \times 10^{-4} \text{ Jy/beam})$ multiplied by $1.8''$ for all integers $n \in [3, 13]$. To improve the visibility of the lobes, a symmetrical logarithmic color scale was used that acts linear in a range within 10^{-2} Jy/beam around 0. *Below:* zoom-ins of the northern and southern lobes respectively. The synthesised beam can be seen in the bottom corners.

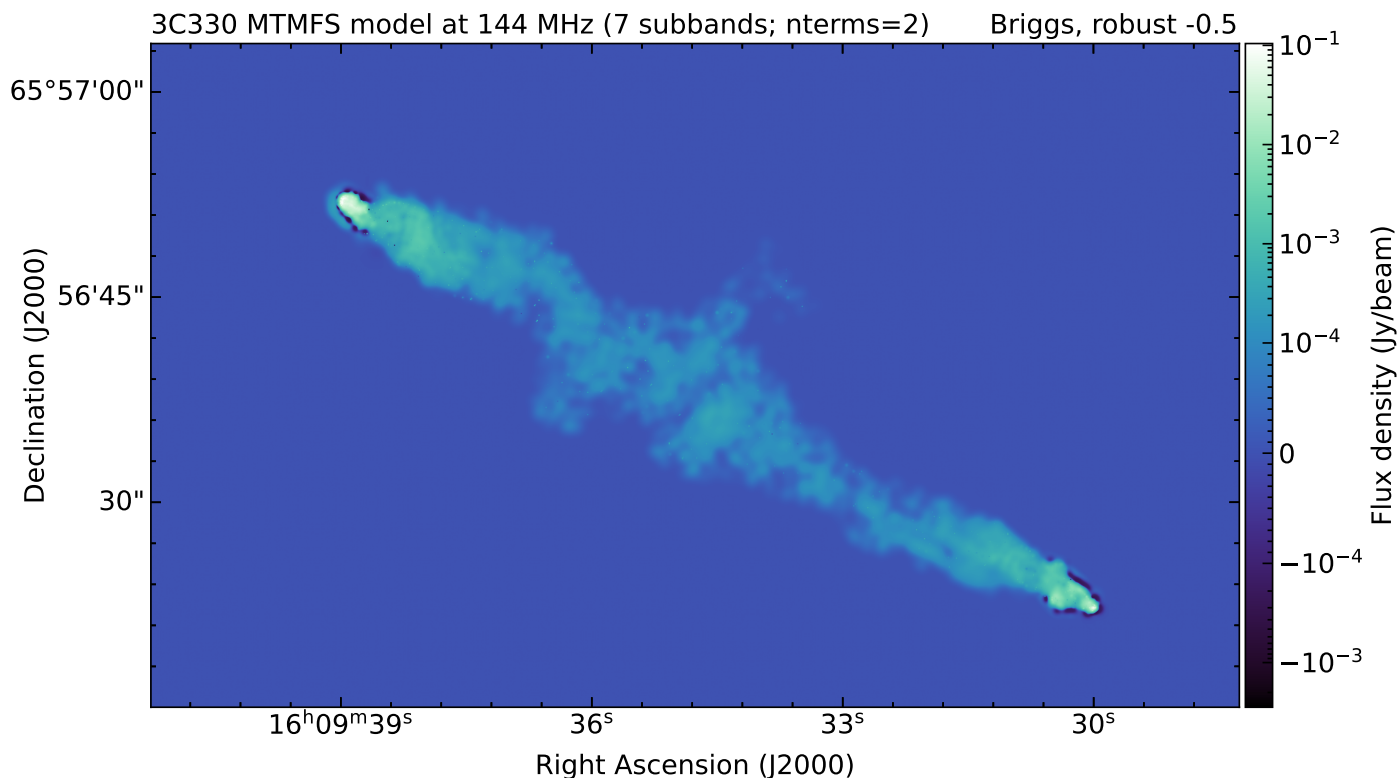


Figure 25: MTMFS model of the seven concatenated subbands. The final model includes scales of 0, 3, 9, 15 and 25 pixels. A symmetrical logarithmic color scale was used that acts linear in a range within 10^{-4} Jy/beam around 0.

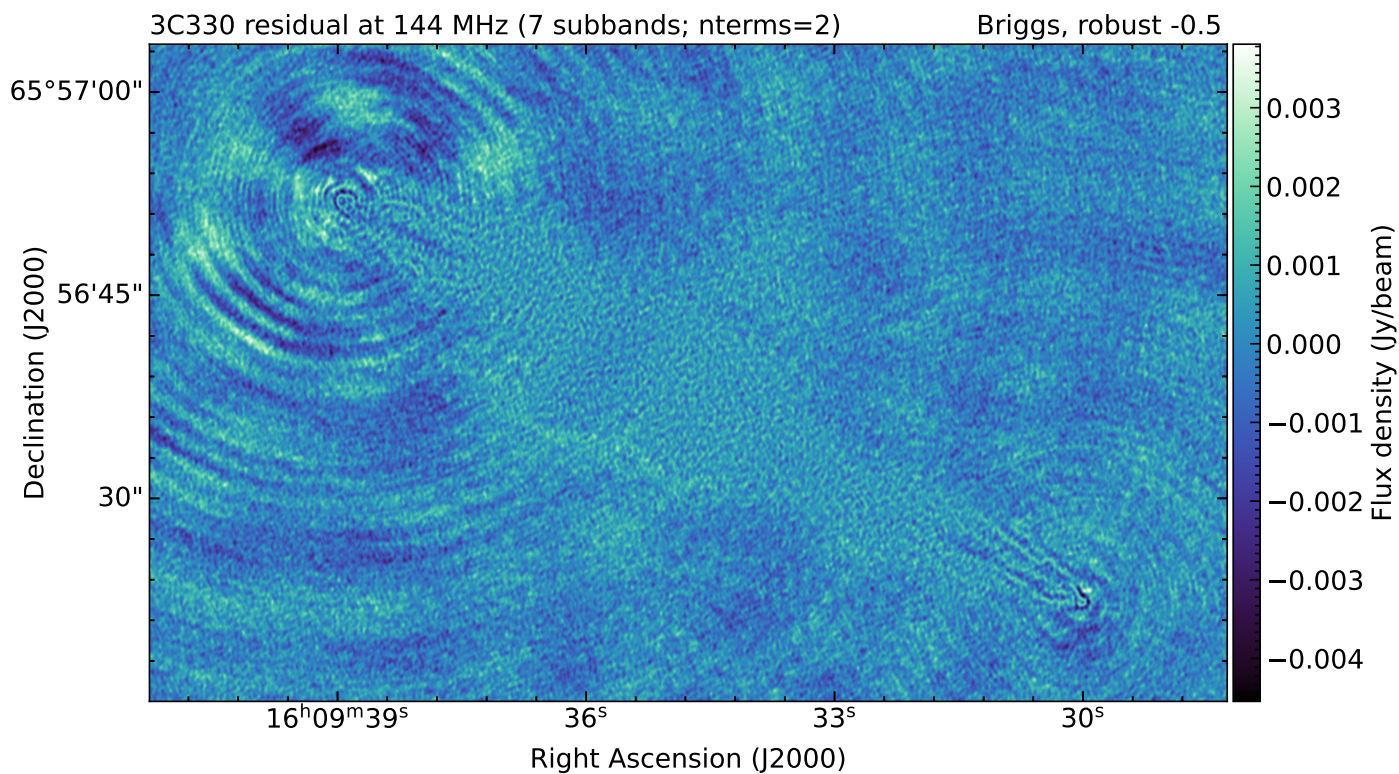


Figure 26: The residuals (data - model) of the seven subband data. Normal scaling is used.




Figure 27: Flux density map of 3C330 at 144 MHz using the fully (averaged) dataset of 260 subbands. Contours are overlaid as $\sigma_{\text{rms},144\text{MHz}} (= 2.932 \times 10^{-4} \text{ Jy/beam})$ multiplied by 1.8^n for all integers $n \in [3, 13]$. To improve the visibility of the lobes, a symmetrical logarithmic color scale was used that acts linear in a range within 10^{-2} Jy/beam around 0.

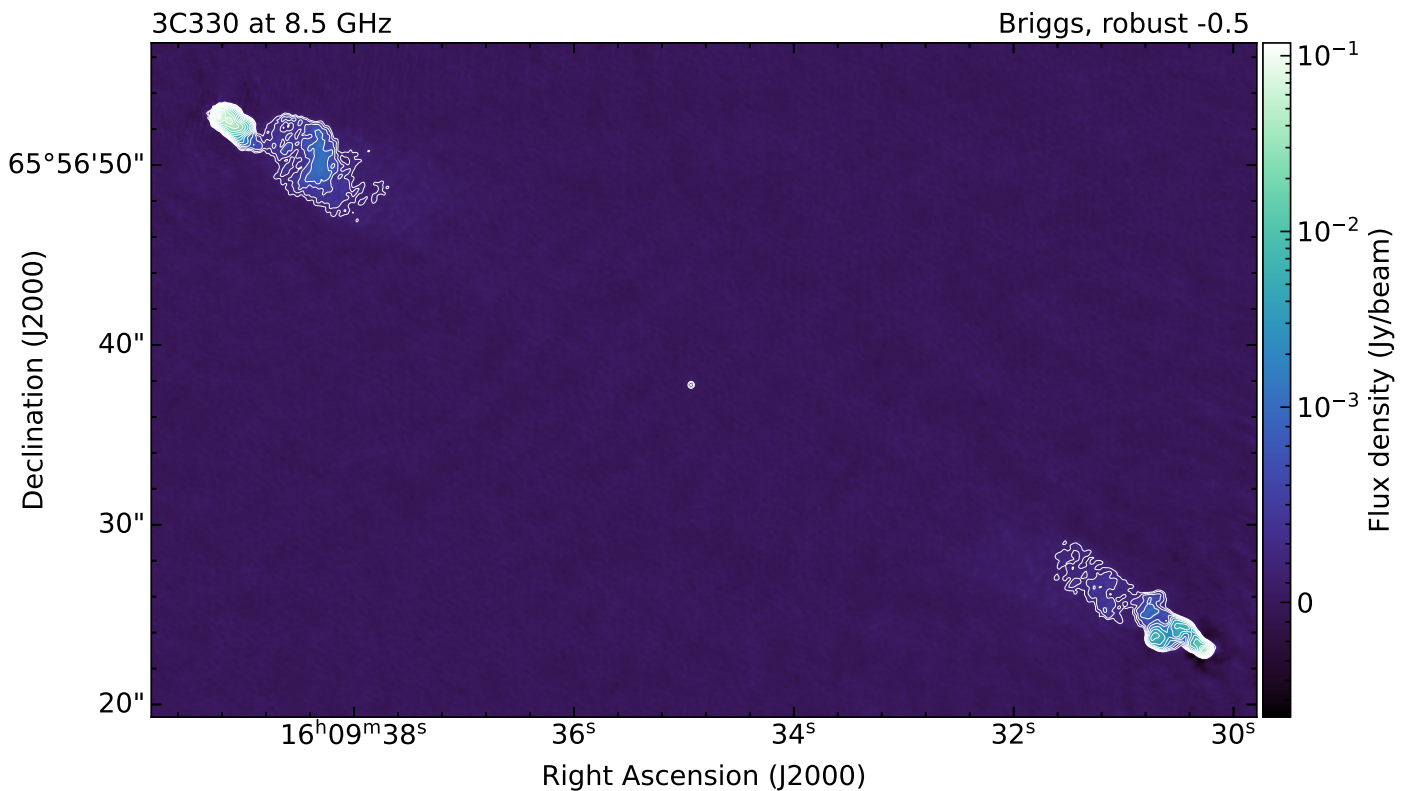


Figure 28: Flux density map of 3C330 at 8.5 GHz. Contours are overlaid as $\sigma_{\text{rms},8.5\text{GHz}} (= 1.880 \times 10^{-5} \text{ Jy/beam})$ multiplied by 1.6^n for all integers $n \in [5, 13]$. To improve the visibility of the lobes, a symmetrical logarithmic color scale was used that acts linear in a range within 10^{-3} Jy/beam around 0. Both lobes and the core can be clearly observed.

5.5 Spectral index map (144 MHz–8.5 GHz)

Our understanding of 3C330’s spectral behaviour can be further analyzed by comparing our data with observations at different frequencies. For this we use the downloaded 8.5 GHz VLA data (see Section 4.6).

The dataset was first imported into a CASA measurement set. After this it was flux calibrated against 3C48 using the corresponding model from the Perley & Butler (2017). Subsequent phase solutions were generated from 1642+690 and all correction tables were applied to the 3C330 data. All of these steps were executed using a simple script. From inspection of the corrected data it was determined that no further manual flagging was required.

After initial corrections, the dataset was split out and subsequently cleaned and self-calibrated in DIFMAP. The applied strategy was to conservatively clean and start with phase-only self-calibration with a solution interval in the order of minutes and slowly lowering that to 10 s. When no improvements in the residual map RMS could be seen, an amplitude selfcal of several minutes was performed. This cycle continued until both phase and amplitude selfcal gave no visible improvements.

For the final image we use WSClean. To ensure good Nyquist sampling we choose a cell size of 0.05'' and convolve the image with a circular Gaussian beam with a diameter of 0.3''. The resulting image of the full source can be seen in Figure 28. We repeat the exact imaging process for our 144 MHz LOFAR data, such that both images have the same cell- and beamsizes.

To align the images we first transform the VLA image to J2000 coordinates using the CASA task `imregrid()`. Afterwards, we used Python to manually align the two images further. Alignment accuracy could easily be checked by creating a spectral index map and looking for extreme values near the lobe edges. The best overlap was found by first aligning on the brightest pixel and subsequently adding an additional shift of 2 pixels in both X and Y.

After optimal alignment we used BRATS to create a two-point spectral index maps of both lobes, which can be seen in Figure 29. BRATS automatically only fits for regions where the flux is 5 times the rms of the background region. Corresponding errormaps are shown in Figure 30.

Table 3: Integrated flux densities of the final maps for both the 7 subband and 260 subband ILT datasets at 144 MHz. ^{ST001} indicates that the dataset contained the phased-up array ST001.

Component + dataset	$S_{144\text{MHz}}$ (Jy)
Full source	
7 subbands ^{ST001}	34.49 ± 0.06
7 subbands	33.32 ± 0.06
260 subbands	34.37 ± 0.04
Northern lobe	
7 subbands ^{ST001}	13.60 ± 0.64
7 subbands	13.20 ± 0.63
260 subbands	14.73 ± 0.47
Southern lobe	
7 subbands ^{ST001}	5.83 ± 0.16
7 subbands	5.94 ± 0.16
260 subbands	5.66 ± 0.13

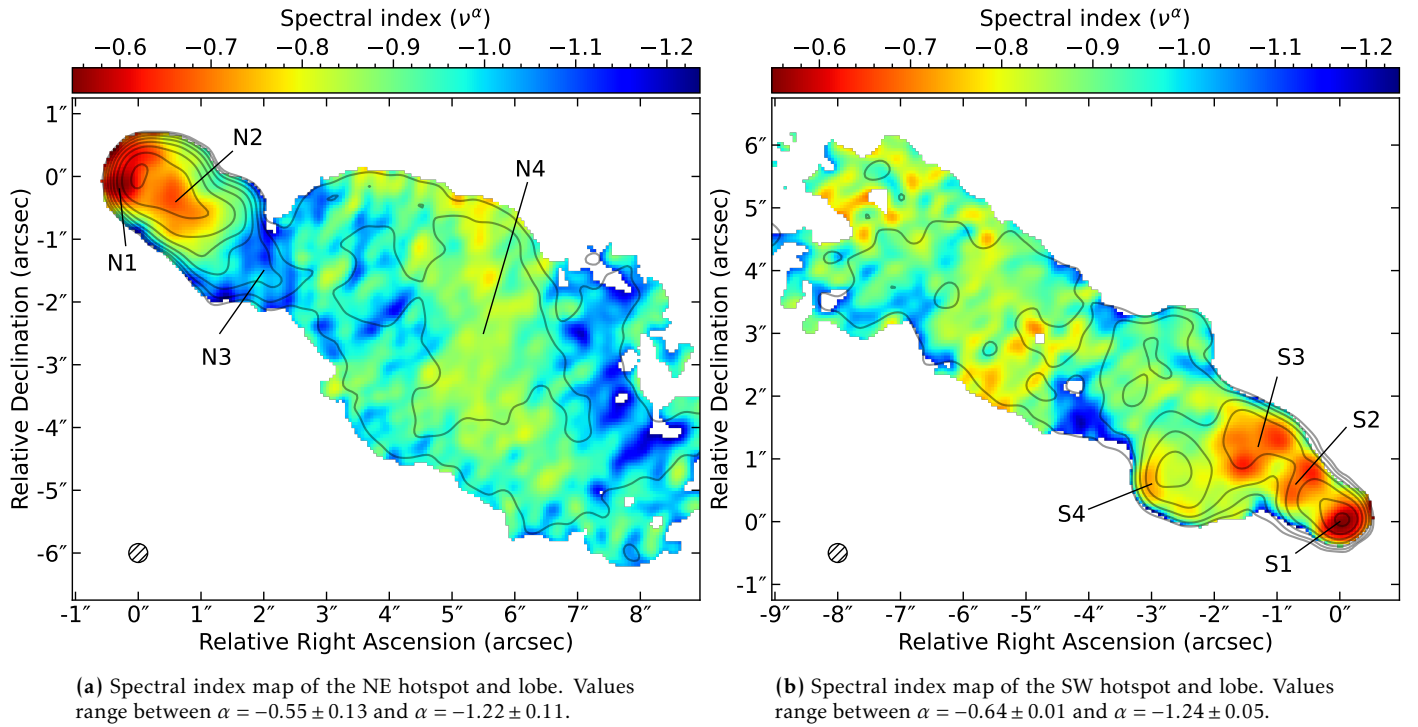


Figure 29: Two-point spectral index map between 144 MHz and 8.4 GHz provided by BRATS. The convolved Gaussian beams size of $0.3'' \times 0.3''$ can be seen in the lower left corner. Contours of the 144 MHz map are overlaid as $\sigma_{\text{rms}, 144 \text{ MHz}} (= 1.233 \times 10^{-3} \text{ Jy/beam})$ multiplied by 2^n for all integers $n \in [3, 15]$. Origin in each map corresponds to the maximum pixel of the 144 MHz flux map. To increase contrast in the colors, the range of values in each map was restricted to $[\alpha_{\text{min}}, 0.9\alpha_{\text{max}}]$.

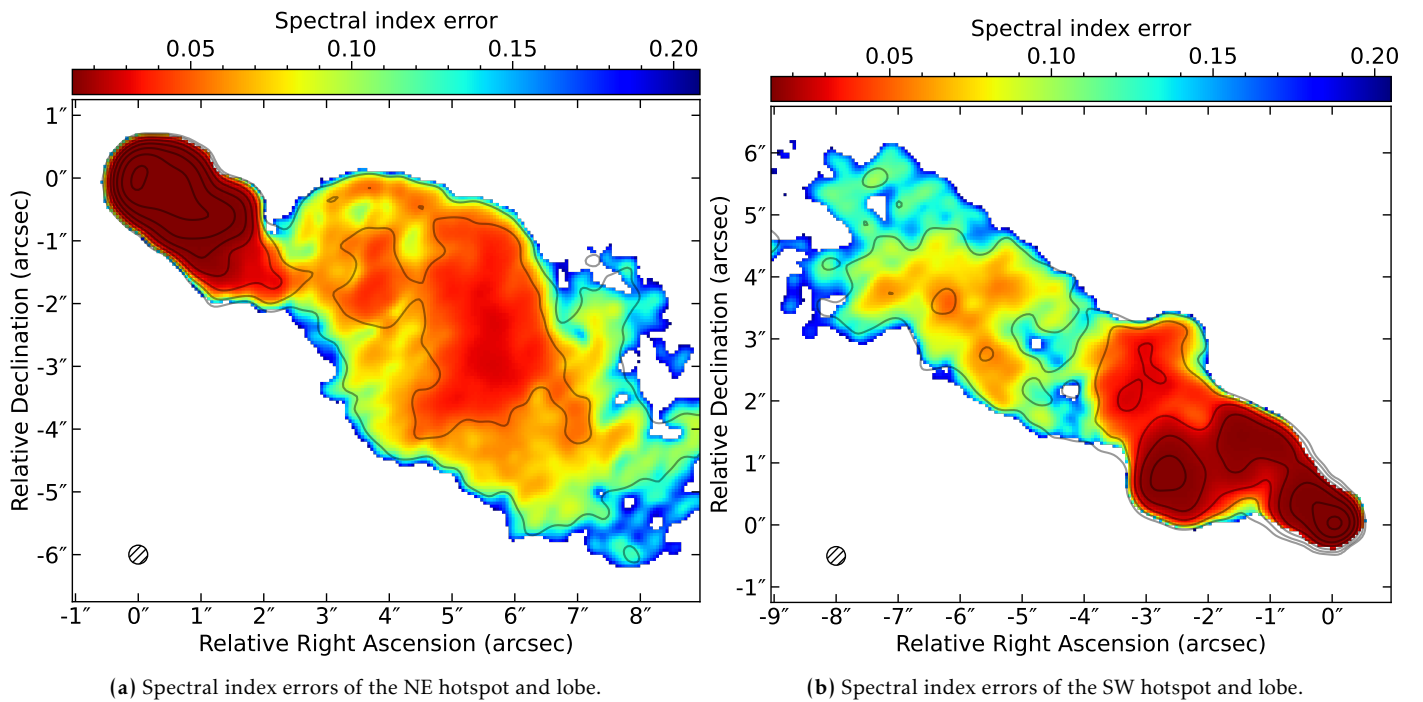


Figure 30: Corresponding fit errors of the two-point spectral index map between 144 MHz and 8.4 GHz provided by BRATS. The convolved Gaussian beams size of $0.3'' \times 0.3''$ can be seen in the lower left corner. Contours of the 144 MHz map are overlaid as $\sigma_{\text{rms}, 144 \text{ MHz}} (= 1.233 \times 10^{-3} \text{ Jy/beam})$ multiplied by 2^n for all integers $n \in [3, 15]$. Origin in each map corresponds to the maximum pixel of the 144 MHz flux map.

6 Discussion

6.1 Surface brightness maps

The obtained surface brightness maps of 3C330 at 144 MHz show a clear detection of both the northern and southern lobe of the source, with sufficient proof of diffuse structure tapering inwards to the core.

The surface brightness map of our 260 subband data is not quite ideal, with the main problem being the large 1'' wide ring-like sidelobes that emanate from the hotspots with a period of around 2''. The northern lobe is mostly affected, as the sidelobe amplitude scales with the emission brightness. The issue with cleaning such data is that the sidelobe emission crosses the diffuse emission just behind both hotspot lobes, and thus will be added to the model as real source structure.

Due to the mentioned issues above and the similarity in results between the 7 subband model and the 260 subband model, we will continue our analysis using the former map. From hereon out all references to the flux density map will thus refer to our 7 subband image (see Figure 24).

Comparison with other intensity maps: Our high-resolution surface brightness maps at 144 MHz closely resemble the previously observed high-resolution maps of [Fernini et al. \(1997\)](#) at 5 GHz, [Gilbert et al. \(2004\)](#) at 8.5 GHz (see Figure 2) and our own 8.5 GHz map, reconstructed from the same observation (see Figure 28).

In the northern lobe we identify a compact primary hotspot with the same size as the synthesised beam. The hotspot is located at the lobe edge and correlates with the location in the 8.5 GHz map. We also recognize the widening of the lobe further towards the nucleus, paired with a local increase of the surface brightness.

The primary hotspot in the southern lobe is compact and located at the furthest edge from the nucleus. We also recognize two regions of increased surface brightness that agree with the locations in the 8.5 GHz map.

Because the Dutch array provides the ILT with an increased sensitivity to the larger scales, we observe more of the diffuse emission that connects the two lobes. We see an increase in surface brightness around the position of where the core should be, but interestingly, we do not detect the core itself. As we are probing the older radiation, the most likely explanation for this is that the electrons in the core have already passed beyond their synchrotron lifetime and thus radiated away all of their energy.

6.2 Spectral index map (144 MHz – 8.5 GHz)

The spectral index maps obtained from BRATS shows complex structure in both the northern and the southern lobe. We will discuss them here.

6.2.1 Spectral structure

The spectral index map of the northern lobe can be divided up into four regions of interest. The flattest part of

the map (N1) lies at the edge of the northernmost lobe region and reaches an average spectral index value of $\alpha_{8500}^{144} \approx -0.55 \pm 0.02$. It can be considered as quite compact and of similar size to the synthesised beam. Of interest is that the peak location does not overlap with the location of the hotspot in the flux map, but is instead slightly shifted to the left of it. The region of the extended tail behind the hotspot (N2) shows a slightly less flattened spectrum of $\alpha_{8500}^{144} \approx -0.70 \pm 0.02$.

The overall steepest spectrum (N3) can be found at the end of the hotspot region and coincides with the low signal-to-noise region in the flux map. The spectrum here shows an average overall spectral index of $\alpha_{8500}^{144} \approx -1.05 \pm 0.05$. The relative flattening of the most eastern region (N4) coincides with an increase of the flux brightness in this area. The spectrum flattens to an average of $\alpha_{8500}^{144} \approx -0.85 \pm 0.04$, and thus contrasts the immediate surrounding steeper parts of the map.

The spectral index map of the southern lobe shows an additional four complex regions of interest. The flattest part (S1) is quite compact and lies at the edge of the southernmost lobe region. This region coincides with the location of the primary hotspot. Similar to the northern lobe, the peak flatness here is about $\alpha_{8500}^{144} \approx -0.64 \pm 0.02$.

In the region behind the primary hotspot we find three more components (S2, S3, S4) with an increased flatness. These components also coincide with an increased surface brightness in the intensity map. Directly in line behind S1 we can observe the components S2 and S3 with a flattened spectral index of $\alpha_{8500}^{144} \approx -0.71 \pm 0.02$. Due to the alignment of S1, S2 and S3, we argue that it is most likely that all three components lie on the jet axis together with the brightest component (N1) of the northern lobe. In both S2 and S3 we see what looks like a double feature, which might indicate a splitting of the jet. Most likely however, is that this is an artifact created due to non-optimal alignment.

A final interesting component is S4. This component is quite compact and lies at the edge of the increased 144 MHz surface brightness contours. It has a flattened spectral index of $\alpha_{8500}^{144} \approx -0.77 \pm 0.02$. Curiously, this component is located furthest away from the proposed jet axis, indicating a potential change in jet direction. By being the steepest of the secondary hotspot components, S4 is also the oldest one.

Our found α_{8500}^{144} spectral index values seem realistic when compared to the standard values obtained from comparable studies ([Mahatma et al. 2023](#), [Kukreti et al. 2022](#)). Our maps also agree with the flattening to $\alpha \sim 0.6$ in the outer lobes of the α_{1418}^{151} map from [Leahy et al. \(1989\)](#) (see Figure 3).

In both lobes a clear primary hotspot undergoing particle acceleration can be identified, both times located at the furthest edge of the observable emission away from the nucleus. We also note the presence of secondary hotspot regions, especially in the southern lobe, and observe that these steepen with distance from the primary hotspot. We will discuss the implications on the dynamics and ori-

gins of these secondary hotspots in Section 6.2.2.

When analysing the spectral index maps we should keep in mind that the process in which they are created is quite sensitive to introducing artifacts. This holds primarily for the edge regions, which easily react to slight errors in alignment by jumping to extreme, non-physical values. From the BRATS errormaps in Figure 30 we conclude that for our science goals a sufficient alignment is achieved. The alignment of the high-resolution VLA and ILT maps can be further improved by for example creating maps with smaller pixel scales (without oversampling) or using weighting schemes that prioritize source accuracy like super-uniform.

6.2.2 Jet acceleration/termination models

There are multiple theories on jet termination and the formation of secondary hotspots. Theories like the ‘dentist-drill’ model (Scheuer 1982) or jet precession models (Cox et al. 1991) suggest that the location of the end-point of the jet changes over time, leaving behind secondary hotspot remnants at the places that were once fuelled by its energetic plasma. As the supply of accelerated particles has ceased, one expects a high-frequency steepening of the spectral index at the locations of these secondary hotspots. These models thus assume a single, arbitrarily long, feeding event of each hotspot after which it is disconnected from the energy supply.

Competing with this theory is the ‘splatter-spot’ model of Williams & Gull (1985), which theorises that the energy from the most compact (primary) hotspot can asymmetrically deflect incoming jet material towards the nearby, more diffuse, (secondary) hotspots. As this is a continuous process, all hotspots stay connected to the initial energy supply of the jet and particle acceleration in the secondary hotspots will not cease after the initial feeding event. This also means that in contrast to the ‘dentist-drill’ model, the spectral index values of these secondary hotspots will stay relatively flat.

The secondary hotspots in 3C330 show a steeper spectral index compared to the primary hotspot, indicating that the plasma in these areas is more aged and that the supply of energy here has been ceased. This observation provides evidence that the most likely acting models in 3C330 are those of the dentist-drill and/or jet precession.

The latter model could furthermore explain the location of the S4 component in the southern lobe, which lies offset from the current jet direction. As the oldest secondary hotspot, it can be argued that a previous state of the jet was about 2° more clockwise in projection. Besides an alignment of the jet axis with S4, this change in precession angle would also mean that the termination point of the jet in the northern lobe is angled more to the top of the bulging N4 region. While there is no indication of a remnant hotspot here, this could give an explanation for the arching structure that defines the northern edge of this region. Moreover, it would agree with the analysis by Gilbert et al. (2004) that the emission structure here indicates a sharp bend (see N1 in Figure 2). It should be noted however, that

this angle is smaller than the inferred 6° of hotspot separation in 3C sources by Best et al. (1995).

From our observations the splatter-spot model is quite unlikely to apply to 3C330, with the main argument being the observed steepening of the secondary hotspots relative to the primary one in our spectral index maps. Moreover, this model would require an extreme jet deflection in the southern lobe of about $\sim 90^\circ$ to account for the existence of the S4 component.

6.3 Issues during post-pipeline cleaning and calibration

Most of our time was spent on the manual calibration of the pipeline corrected data. In this time numerous issues were encountered, and various solutions were attempted to solve them. The most impactful issues will be discussed here:

MT-MFS cleaning: The most commonly encountered issue with MT-MFS cleaning has been the formation of negative bowls around the bright hotspots of the source. This was encountered both in CASA and in WSCLEAN and occurred due to the use of large scales (more than twice the size of the current brightest component in the mask) when cleaning the hotspots. The negative bowls can be clearly seen in Figure 25. It was found that this issue could be remedied somewhat by using a very restrictive mask during the cleaning process, inking towards a few pixels inwards from the source edge. The best found method however, was to use a variable mask with a customised subset of scales that best describe its flux content. In this way, the hotspot lobe should be cleaned using a small mask together with more compact scales (e.g. 0, 3 pixels), and the diffuse middle section should be cleaned using a mask that excludes the hotspot region and fits for larger scales (e.g. 9, 12, 15 pixels).

XX and YY not aligned: During cleaning we found that there was a slight offset between the XX and YY flux maps, leading to large amplitude sidelobes around the hotspot areas. Additionally, the YY polarisation maps showed a large amount of offset flux to the left of the NE hotspot. To remedy this, we continued using only the XX polarisation in calibration beyond the seven subband model. Doing this decreases our visibilities by a factor of 2 and consequently increases the rms noise by a factor of $\sqrt{2}$. We can justify doing this, as our source is very bright and we have sufficient S/N to still model accurately. Besides fixing the sidelobe problem, this also increased computation time.

Removing ST001: Through initial inspection of the data it was found that the NE hotspot was much higher in flux density than the SW hotspot. This is due to a combination of the following two facts: the phase centre of the dataset lies closer to the NE hotspot ($\sim 20''$) than to the SW hotspot ($\sim 35''$) and the data gets attenuated by the large beam of the combined station ST001 ($\theta_{\text{FWHM}} \approx 100''$). Due to it’s

origin, ST001's beamsize makes a large contribution to the total synthesised beam of the observation and, in addition, is the station with by far the largest sensitivity. This all contributes to the fact that the SW hotspot gets attenuated relatively more than the NE hotspot, creating an artificial flux imbalance. By splitting out ST001 from our dataset we remove the additional attenuation of the SW hotspot and we get a more accurate flux balance between the two hotspots.

From Table 3 we see that the ratio of the integrated flux densities goes from $S_{\text{NE}}/S_{\text{SW}} = 2.33 \pm 0.12$ to $S_{\text{NE}}/S_{\text{SW}} = 2.22 \pm 0.12$ when we remove the station ST001. While the difference is subtle, it does contribute to an imbalance in the total fluxes between these hotspots.

Besides a change in the overall flux distribution, removing the phased-up array did not alter any of the structure seen in the bright lobes. This can be attributed to the fact that ST001 was mainly probing the diffuse structures in-between the core and lobes. Consequently, this did mean that the removal of ST001 slightly lowered our sensitivity in these areas, but as it was already noticed early on that high-resolution spectral map imaging in these regions was not realistic, this did not form a noticeable problem.

Dominant sidelobes in the residuals: For the calibration of the individual subbands, even after using only the XX polarisation, we still found large, wide sidelobes propagating from the hotspot origins. These features can be observed most notably originating from the NE hotspot, but can also be seen originating from the SW hotspot. A side-lobe pattern with a fringe width of about $\sim 2''$ (measured from peak to peak in the CASA task `imview`) forms the main problem in our imaging. We can use this fringe width measurement to try and deduce which baseline could be responsible. Using the frequency range of the observation (121 MHz – 166 MHz) and the fact that:

$$B = \frac{\lambda}{W_{\text{fringe}}}, \quad (6.1)$$

where B is our baseline length, λ is our observing wavelength and W_{fringe} is our fringe width, we can say that our troublesome baselines lie in the range between:

$$B_{121\text{MHz}} \approx 255 \text{ km} \quad \text{and} \quad B_{166\text{MHz}} \approx 186 \text{ km}.$$

From baseline inspection we see that there are about 30 baselines in this range and most of these baselines run between one of the Dutch remote stations (RS) and one of the German stations (most notably DE605 and DE609). In an attempt to fix this problem the potentially problematic baselines and stations were inspected using `plotms` and adequately flagged where needed. Besides some time bin outliers, these baselines did however not exhibit any other weird features and their visibility amplitudes were in line with the rest of the observation.

7 Conclusion

In this thesis we calibrated new high angular resolution ILT data in the low frequency regime (144 MHz) to analyse the spectral properties of bright radio galaxy 3C330. We initially used the LOFAR-VLBI pipeline to fully calibrate the Dutch array and to calibrate DDEs in the international stations. As our source gave problems during the calibration of DDEs, we continued with a manual calibration outside of the pipeline.

By utilizing a strategy of manually calibrating seven subbands, we managed to obtain the first $0.3''$ resolution map of 3C330 at 144 MHz and by aligning this with our calibrated data of archival VLA data at 8.5 GHz, we have also created the first high-resolution spectral index map of the 3C330 lobes. A complex spectral structure is observed in both lobes and no detection of the core is made.

The detected hotspots in the lobes show proof of both recent and older activity of accelerated particles. Our analysis of the spectral structure suggests that the evolution of the jet interaction with the surrounding medium can best be described by a combination of the jet precession and 'dentist-drill' models, indicating that only the primary hotspot is currently being fuelled by the energetic jet material. We suggest that the off-axis S4 component of the source can be explained by a previous evolution of the jet precession angle of about 2° , which is smaller than the inferred precession angles of 3C sources. Overall, the source 3C330 acts as a typical FR-II radio galaxy that can be well-explained by the current models for radio galaxy evolution.

From our study it can be seen that the ability to spatially resolve the properties of both the small and the large scale emission in the low-frequency radio regime can be of great importance in further understanding the evolution of radio galaxy objects. The unique combination of high resolving power of the small scales and the high sensitivity on the large scales makes the ILT an excellent instrument for this job. It will be exciting to see future radio galaxy studies that come from the calibration of the incredibly large amount of data that the ILT has already provided (LoTSS DR2) and will continue to provide in the future.

7.1 Technical issues and future recommendations

Using the LOFAR-VLBI pipeline, most of the calibration could be done relatively successfully. However, like most VLBI data, the calibration of the ILT data remains quite tricky and prone to error, especially with large spatial sources like 3C330. We find that in such cases the pipeline could not create an adequate model for direction-dependent delay calibration and thus finds inadequate TEC solutions for the field.

In this thesis we show (semi-)successful proof of a manual calibration strategy that continues on the calibration done in the LOFAR-VLBI pipeline. While this strategy is quite specific to our case, it might prove useful for similar cases in which insufficient TEC solutions of the dis-

persive delay calibrator are found. Here we note some of our recommendations and improvements for this process:

- The pipeline provides limited information on the calibrated visibilities of the data. We therefore highly suggest a check and flagging of the visibility data with a tool like CASA before continuing on with the manual calibration.
- During the MT-MFS cleaning of very bright sources, use a tightly defined mask to limit the addition of large negative model components and therefore negative bowl forming around the source.
- Identify and remove the phased-up core ST001 during cleaning and calibration if the source has a large angular size or lies far ($\sim 30''$) from the phase centre.

Besides technical recommendations, there are a couple of ways that this research can be extended upon.

Firstly, we note that by using only seven subbands for our final results, we have not used the full potential of our observation. While the reasons for our data issues remain partly unclear, these problems might get resolved by attempting different calibration and/or flagging strategies. A sufficient calibration of the full, unaveraged dataset could provide us with the missing information on the diffuse structures inbetween the two lobes.

Secondly, information like spectral curvature of 3C330 could be obtained by adding additional high-resolution observations at other frequencies. The foremost candidates for this are the EVLA observations at 12 and 15 GHz, currently available on the NRAO Data Archive⁹.

Finally, the calibrated model of 3C330 could be used to calibrate the data field P241+65 for DfEs. This field contains e.g. a quadruply lensed quasar (B1609+656), which, when mapped in sub-arcsecond resolution, could become an interesting asset to the field of gravitational-lens modelling.

8 Acknowledgements

As an expression of gratitude I hereby present a brief list of people that have helped me complete this thesis.

Firstly, the LOFAR long baseline group for support and help with the problems that arose with running the pipeline, assessing the quality solutions and being available for any questions. In particular I would like to thank Leah Morabito for her help with transitioning from within the pipeline to work outside of it.

I'd like to thank Timothy Shimwell for the help with staging the data from the LTA, setting up the pipeline and providing the scripts for downloading and extracting the LTA data and Pranav Kukreti for helping me running the pipeline. Likewise, I'd like to thank Emilio Ceccotti for answering any questions that I had on using WSCLEAN.

I would like to thank the various people at Kapteyn, both inside and outside the radio astronomy group, for their

general support, interesting conversations and friendly atmosphere that they provided during the entire process. Particular shout-outs go to Ramon Luichies, Jelke Bethlehem, Jasper Steringa, Bram Alferink and Celestin Herbe-George.

Lastly and most certainly not least, I'd like to extend my special thanks to my supervisor, John McKean, for his continued support during the process of completing my thesis and seeing it through with me.

⁹<https://data.nrao.edu/portal/#/>

References

- Artyukh V. S., Dagkesamanskii V. V., Vitkevich R. D., Kozhukhov V. N., 1969, *Soviet Astronomy*, **12**, 567
- Baade W., Minkowski R., 1954, *(the) Astrophysical Journal*, **119**, 206
- Baars J. W. M., Genzel R., Pauliny-Toth I. I. K., Witzel A., 1977, *A&A*, **61**, 99
- Barthel P. D., 1989, *(the) Astrophysical Journal*, **336**, 606
- Bennett A. S., 1962, *Monthly Notices of the RAS*, **125**, 75
- Best P. N., Bailer D. M., Longair M. S., Riley J. M., 1995, *Monthly Notices of the RAS*, **275**, 1171
- Blandford R. D., Rees M. J., 1974, *Monthly Notices of the RAS*, **169**, 395
- Bridle A. H., Schwab F. R., 1999, in Taylor G. B., Carilli C. L., Perley R. A., eds, *Astronomical Society of the Pacific Conference Series Vol. 180, Synthesis Imaging in Radio Astronomy II*. p. 371
- Cohen A. M., Porcas R. W., Browne I. W. A., Daintree E. J., Walsh D., 1977, *Memoirs of the RAS*, **84**, 1
- Cohen A. S., Lane W. M., Cotton W. D., Kassim N. E., Lazio T. J. W., Perley R. A., Condon J. J., Erickson W. C., 2007, *Astronomical Journal*, **134**, 1245
- Cox C. I., Gull S. F., Scheuer P. A. G., 1991, *Monthly Notices of the RAS*, **252**, 558
- Douglas J. N., Bash F. N., Bozyan F. A., Torrence G. W., Wolfe C., 1996, *Astronomical Journal*, **111**, 1945
- Edge D. O., Shakeshaft J. R., McAdam W. B., Baldwin J. E., Archer S., 1959, *Memoirs of the RAS*, **68**, 37
- English W., Hardcastle M. J., Krause M. G. H., 2016, *Monthly Notices of the RAS*, **461**, 2025
- Fanaroff B. L., Riley J. M., 1974, *Monthly Notices of the Royal Astronomical Society*, **167**, 31P
- Fernini I., 2002, *Astronomical Journal*, **123**, 132
- Fernini I., Burns J. O., Perley R. A., 1997, *Astronomical Journal*, **114**, 2292
- Frey S., Mosoni L., 2009, *New Astronomy Reviews*, **53**, 307
- Genzel R., Pauliny-Toth I. I. K., Preuss E., Witzel A., 1976, *Astronomical Journal*, **81**, 1084
- Gilbert G. M., Riley J. M., Hardcastle M. J., Croston J. H., Pooley G. G., Alexander P., 2004, *Monthly Notices of the RAS*, **351**, 845
- Gregory P. C., Condon J. J., 1991, *Astrophysical Journal, Supplement*, **75**, 1011
- Hales S. E. G., Masson C. R., Warner P. J., Baldwin J. E., 1990, *Monthly Notices of the RAS*, **246**, 256
- Hardcastle M. J., 2018, *Monthly Notices of the RAS*, **475**, 2768
- Hardcastle M. J., Alexander P., Pooley G. G., Riley J. M., 1997, *Monthly Notices of the RAS*, **288**, 859
- Hardcastle M. J., Birkinshaw M., Cameron R. A., Harris D. E., Looney L. W., Worrall D. M., 2002, *The Astrophysical Journal*, **581**, 948
- Högbom J. A., 1974, *Astronomy and Astrophysics, Supplement*, **15**, 417
- Intema H. T., 2009, PhD thesis, Leiden Observatory
- Intema H. T., Jagannathan P., Mooley K. P., Frail D. A., 2017, *Astronomy & Astrophysics*, **598**, A78
- Jackson et al., 2016, *A&A*, **595**, A86
- Jackson et al., 2022, *A&A*, **658**, A2
- Kellermann K. I., Pauliny-Toth I. I. K., Williams P. J. S., 1969, *(the) Astrophysical Journal*, **157**, 1
- Kukreti P., et al., 2022, *A&A*, **658**, A6
- Laing R. A., 1981, *Monthly Notices of the RAS*, **195**, 261
- Laing R. A., Peacock J. A., 1980, *Monthly Notices of the Royal Astronomical Society*, **190**, 903
- Laing R. A., Riley J. M., Longair M. S., 1983, *Monthly Notices of the RAS*, **204**, 151
- Leahy J. P., Muxlow T. W. B., Stephens P. W., 1989, *Monthly Notices of the RAS*, **239**, 401
- Leahy J. P., Black A. R. S., Dennett-Thorpe J., Hardcastle M. J., Komissarov S., Perley R. A., Riley J. M., Scheuer P. A. G., 1997, *Monthly Notices of the RAS*, **291**, 20
- Mahatma V. H., Basu A., Hardcastle M. J., Morabito L. K., van Weeren R. J., 2023, *Monthly Notices of the RAS*, **520**, 4427
- McCarthy P. J., van Breugel W., Kapahi V. K., 1991, *(the) Astrophysical Journal*, **371**, 478
- McKean J. P., et al., 2016, *Monthly Notices of the Royal Astronomical Society*, **463**, 3143
- Montenegro-Montes F., 2009, PhD thesis
- Morabito L. K., et al., 2022, *Astronomy & Astrophysics*, **658**, A1
- Offringa A. R., 2016, *A&A*, **595**, A99
- Offringa A. R., van de Gronde J. J., Roerdink J. B. T. M., 2012, *Astronomy & Astrophysics*, **539**, A95
- Pauliny-Toth I. I. K., Witzel A., Preuss E., Kühr H., Kellermann K. I., Fomalont E. B., Davis M. M., 1978, *Astronomical Journal*, **83**, 451
- Perley R. A., Butler B. J., 2017, *The Astrophysical Journal*, **230**, 7
- Perley R. A., Dreher J. W., Cowan J. J., 1984, *(the) Astrophysical Journal Letters*, **285**, L35
- Perucho M., López-Miralles J., 2023, *arXiv e-prints*, p. [arXiv:2306.05864](https://arxiv.org/abs/2306.05864)
- Pyrzas S., Steenbrugge K. C., Blundell K. M., 2015, *A&A*, **574**, A30
- Rees N., 1990a, *Monthly Notices of the RAS*, **243**, 637
- Rees N., 1990b, *Monthly Notices of the RAS*, **244**, 233
- Rengelink R. B., Tang Y., de Bruyn A. G., Miley G. K., Bremer M. N., Roettgering H. J. A., Bremer M. A. R., 1997, *Astronomy and Astrophysics, Supplement*, **124**, 259
- Roger R. S., Costain C. H., Bridle A. H., 1973, *Astronomical Journal*, **78**, 1030
- Roger R. S., Costain C. H., Stewart D. I., 1986, *Astronomy and Astrophysics, Supplement*, **65**, 485
- Rybicki G. B., Lightman A. P., 1979, *Radiative processes in astrophysics*
- Scaife A. M. M., Heald G. H., 2012, *Monthly Notices of the Royal Astronomical Society: Letters*, **423**, L30
- Scheuer P. A. G., 1974, *Monthly Notices of the RAS*, **166**, 513
- Scheuer P. A. G., 1982, in Heeschen D. S., Wade C. M., eds, *Vol. 97, Extragalactic Radio Sources*. pp 163–165

- Shepherd M. C., 1997, in Hunt G., Payne H., eds, *Astronomical Society of the Pacific Conference Series Vol. 125, Astronomical Data Analysis Software and Systems VI*, p. 77
- Shimwell T. W., et al., 2017, *Astronomy & Astrophysics*, 598, A104
- Shimwell T. W., et al., 2019, *Astronomy & Astrophysics*, 622, A1
- Shimwell T. W., et al., 2022, *Astronomy&Astrophysics*, 659, A1
- Singal A. K., 2014, *Astronomical Journal*, 148, 16
- SingularityCE D., 2021, SingularityCE 3.8.3, doi:10.5281/zenodo.5564915, <https://doi.org/10.5281/zenodo.5564915>
- Smirnov O. M., 2011, *A&A*, 527, A107
- Stappers, B. W. et al., 2011, *A&A*, 530, A80
- Steenbrugge K. C., Blundell K. M., Duffy P., 2008, *Monthly Notices of the RAS*, 388, 1465
- Thompson A., Moran J. M., Swenson Jr G. W., 2017, *Interferometry and synthesis in radio astronomy*. Springer Nature
- Timmerman, R. et al., 2022, *A&A*, 658, A5
- Turner R. J., Shabala S. S., 2015, *(the) Astrophysical Journal*, 806, 59
- Williams A. G., Gull S. F., 1985, *Nature*, 313, 34
- Williams W. L., et al., 2016b, *Monthly Notices of the Royal Astronomical Society*, 460, 2385
- Williams W. L., et al., 2016a, *Monthly Notices of the Royal Astronomical Society*, 460, 2385
- Yates-Jones P. M., Turner R. J., Shabala S. S., Krause M. G. H., 2022, *Monthly Notices of the RAS*, 511, 5225
- de Gasperin F., et al., 2019, *Astronomy & Astrophysics*, 622, A5
- van Haarlem M. P., et al., 2013a, *Astronomy & Astrophysics*, 556, A2
- van Haarlem M. P., et al., 2013b, *Astronomy & Astrophysics*, 556, A2
- van Weeren R. J., et al., 2016, *The Astrophysical Journal, Supplement*, 223, 2

Appendices

A Data inspection

A.1 Diagnostic plots of Pre-Facet-Calibrator.parset

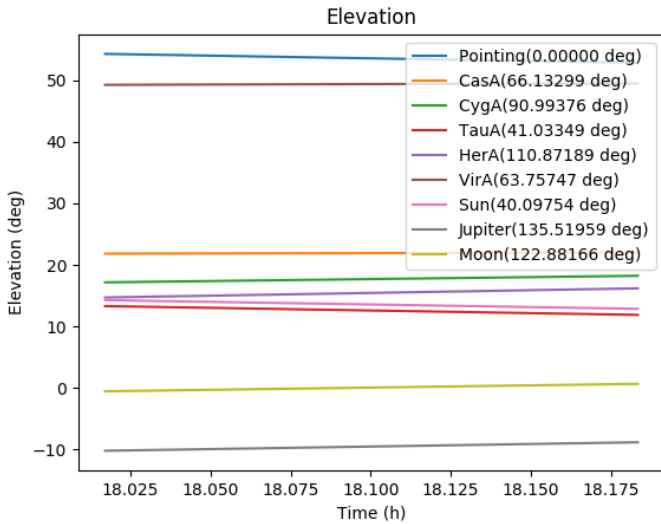


Figure 31: The elevation of A-Team sources (Cassiopeia A, Cygnus A, Taurus A, and Virgo A) together with the brightest celestial bodies during the observation of the standard calibrator 3C196. Labelled in brackets the legend are the distances to the respective sources.




Figure 32: Matrix plot of the derived instrumental clock offsets in seconds. A constant offset is calculated for each station.

A.2 Diagnostic plots of Pre-Facet-Target.parset




Figure 33: The elevation of A-Team sources (Cassiopeia A, Cygnus A, Taurus A, and Virgo A) together with the brightest celestial bodies as a function of time during the observation of the target field P241+65. Labelled in brackets the legend are the distances to the respective sources.




Figure 34: The fraction of flagged data that were flagged due to potential contamination from the A-team sources versus frequency.




Figure 35: The remaining fraction of unflagged data after PRE-FACTOR versus frequency.

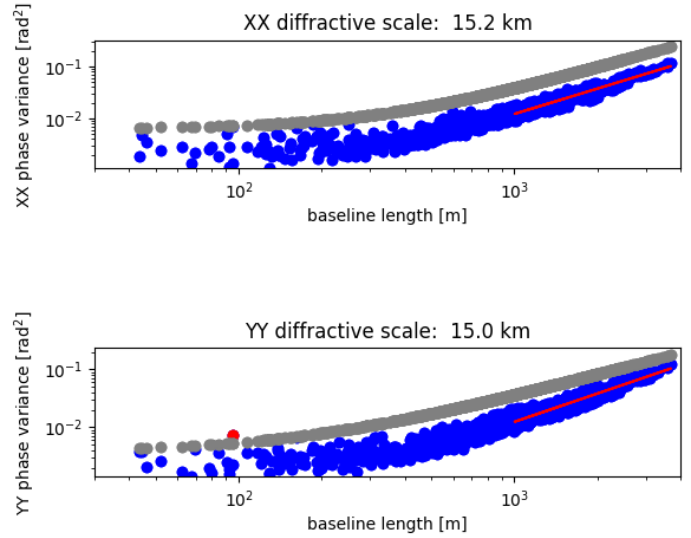


Figure 36: Plots of ionospheric structure of the target field P241+65 as a function of phase variance versus baseline length.

A.3 Diagnostic plots of Delay-Calibration.parset




Figure 37: Pipeline generated image of 3C330 with the resolution of $\sim 6''$ of the Dutch array (i.e. using a $u-v$ cut of $50k\lambda$ to exclude baselines to international stations). The phase centre is positioned at the location of the N-W lobe.

A.4 Diagnostic plots in using CASA's plotms




Figure 38: The visibility amplitudes versus time of a single (averaged) spectral window at 150 MHz. Manually flagged data can be seen in red.

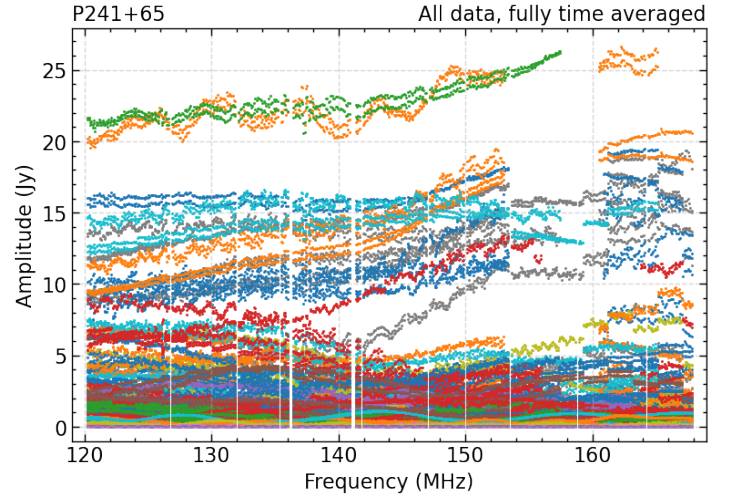


Figure 39: The visibility amplitudes versus frequency. The data has been averaged over the entire observation time and is colorized by station.

B Complementary results

B.1 Sources for the spectral energy distribution (SED)

Table 4: Used literature flux density measurements for the SED of 3C330.

¹Measurements with no given uncertainty were given a 10% error. ²Flux density scales: RCB; Roger et al. (1973), B77; Baars et al. (1977), TXS; Douglas et al. (1996), KPW; Kellermann et al. (1969). ³If a factor is given, the original measurement was mapped to the RCB or KPW scale by application of the corresponding factor taken from ^aScaife & Heald (2012), an interpolation of the factors given in ^bBaars et al. (1977) or a factor from ^cDouglas et al. (1996).

Freq. [MHz]	Flux density [Jy]	Uncertainty ¹	Flux scale ²	Factor ³	Notes	Reference for measurement
10	110	± 60	RCB	–		Roger et al. (1973)
22	91	± 6	RCB	–		Roger et al. (1986)
38	85.50	± 8.55	RCB	–	8C	Rees (1990b)
	95.58	± 9.56	RCB	1.18 ^a		Kellermann et al. (1969)
74	54.07	± 5.47	KPW	1.020 ^b	VLSS	Cohen et al. (2007)
86	57.15	± 4.98	RCB	0.94 ^a		Artyukh et al. (1969)
152	34.59	± 0.090	RCB	–	6C	Hales et al. (1990)
178	30.302	± 1.515	RCB	1.09 ^a		Kellermann et al. (1969)
325	18.81	± 0.76	RCB	0.9 ^a	WENSS	Rengelink et al. (1997)
365	21.379	± 0.272	KPW	0.977 ^c	Texas	Douglas et al. (1996)
408	16.8	–	KPW	–		Singal (2014)
750	10.40	± 0.52	KPW	–		Kellermann et al. (1969)
966	8.72	± 0.69	KPW	–		Cohen et al. (1977)
1400	6.90	± 0.35	KPW	–		Kellermann et al. (1969)
2700	3.76	± 0.19	KPW	–		Kellermann et al. (1969)
4850	2.422	± 0.215	KPW	0.993 ^b		Gregory & Condon (1991)
4900	2.38	± 0.08	KPW	–		Pauliny-Toth et al. (1978)
5000	2.35	± 0.12	KPW	–		Kellermann et al. (1969)
8470	1.2	–	–	–		Gilbert et al. (2004)
10700	1.02	± 0.03	KPW	–		Pauliny-Toth et al. (1978)
14900	0.609	± 0.039	KPW	0.966 ^b		Genzel et al. (1976)

B.2 8.5 GHz VLA observation

Table 5: Observation parameters of the 8.5 GHz VLA dataset: ¹Very Large Array, ²Averaged down from raw data by a factor of 10, ³Given in B1950 right ascension (RA) and declination (DEC).

Dataset properties	
Telescope ¹	VLA
Project name	AP331
Observation date	14 December 1996
Configuration	A
Maximum baseline	35.5 km
Time resolution ²	10 s
Frequency resolution ²	25 MHz
Central frequency	8.475 GHz
Bandwidth	50 MHz
Flux density calibrator	
Calibrator name	3C48
Phase center ³	01 ^h 34 ^m 49.83 ^s +32°54′20.52″
Time on source	2 minutes 30 seconds
Phase calibrator	
Calibrator name	1642+690
Phase center ³	16 ^h 42 ^m 18.06 ^s +69°2′13.22″
Time on source	4 minutes
Target	
Target name	3C330
Phase center ³	16 ^h 09 ^m 13.9 ^s +66°4′22.3″
Time on source	90 minutes

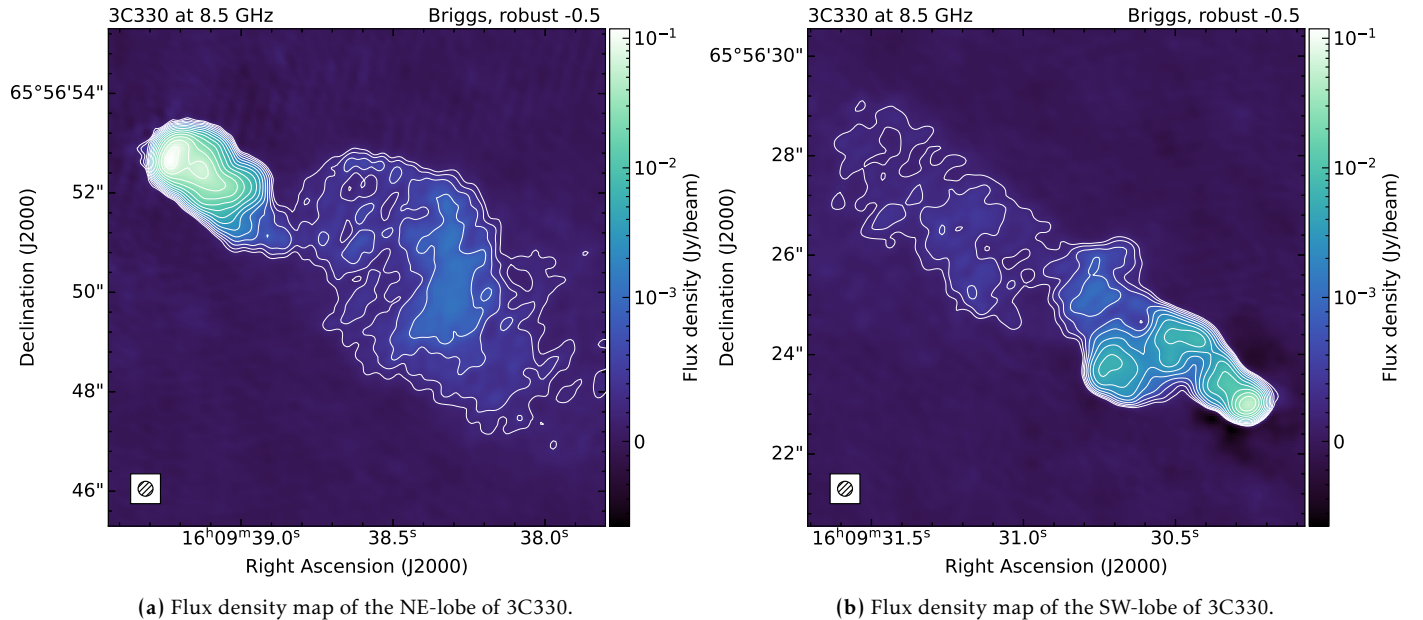


Figure 40: Flux density map of 3C330 at 8.5 GHz. The restoring beam of $0.3'' \times 0.3''$ can be seen in the lower left hand corner. Contours are overlaid as $\sigma_{\text{rms}, 8.5 \text{ GHz}} (= 1.880 \times 10^{-5} \text{ Jy/beam})$ multiplied by 1.6^n for all integers $n \in [5, 30]$. To improve the visibility of the lobes, a symmetrical logarithmic color scale was used that acts linear in a range within 10^{-3} Jy/beam from 0.