UNIVERSITY OF GRONINGEN

MASTER THESIS

Investigating spectral flattening of bright radio features in Cygnus A

Using 110 to 250 MHz LOFAR observations

March 28, 2021

rijksuniversiteit groningen

Author:
J.T. Steringa

Supervisor:
prof. dr. J.P. McKean

Abstract

Cygnus A is the archetype example of an FR-II radio galaxy, being one of the brightest objects in the radio sky. Imaging Cygnus A at radio wavelengths can give valuable insight to the formation of hotspots in radio galaxies of this type, as well as other phenomena that act in creating the complex structure seen in its radio lobes. Observations with the Low-Frequency Array (LOFAR) were used to image Cygnus A between 110 and 250 MHz at an angular resolution of ~ 2 arcseconds. The four known hotspots are seen, and more diffuse structure throughout the lobes. A spectral index and curvature map show that the brightest regions in the lobes have flatter and more strongly curved spectra than generic lobe plasma. Seven regions are chosen for a deeper spectral analysis. A spectral turnover is seen in secondary hotspots A and D, while primary hotspot B only shows flattening towards 110 MHz. A faint blob near the edge of the counter-lobe is proposed as a relic hotspot candidate, as its spectrum shows hotspot-like behaviour. Synchrotron self-absorption and free-free absorption models reasonably fit the LOFAR spectra of the hotspots combined with Karl G. Jansky Very Large Array (VLA) data, but the inferred magnetic field strength and electron density take unphysical values. An extended bright semicircle shaped region and two parallel filaments in the counter-lobe show behaviour consistent with shock wave producing events, as they radiate significantly brighter than generic counter-lobe plasma.

Acknowledgements

The one and a half years of my master research project have been really exciting, as it enabled me to commit to a subject of my passion for a very long time. It was also a two-sided time, as the COVID-19 pandemic struck early 2020, after the first six months of research. I would like to say thanks to a few people who have motivated and supported me in my work, both before and during the pandemic.

Firstly, I would like to thank John McKean for his supervision. John, thank you for all the highly appreciated time and effort you spent guiding me during my project. Finishing the project took a little longer than expected, due to difficulties working from home. Nevertheless, you always stayed patient and willing to help, even when things came later than hoped. This I sincerely appreciate. The daily Zoom meetings with the group were really valuable by getting insight in everyone's research, as well as the much needed personal contact with everyone.

Secondly, I would like to thank everyone in the research group: Ruslan, Samira, Ramon, Chatrik, Arno, Nanna and Di for your kindness and warmth during the daily Zoom meetings. You gave a pleasant and fascinating look into all your projects, and I also really enjoyed the many talks and laughs about our daily activities.

Thirdly, I would like to thank all my fellow astronomy students with whom I shared the student office (or as some of us called it, the "therapy room") during working hours at the Kapteyn Institute, before the pandemic abruptly ended our daily routine. Sharing thoughts on everyone's findings or frustrations, getting on the coffee train, and of course fooling around by throwing paper airplanes and doing somersaults were the kinds of distractions that I needed every now and then, and it kept the room alive. I really hope to see you all again when things get back to normal, I will really miss the atmosphere during our days in the office.

In addition, I would like to express my appreciation to Martin, Leon and Eite of the Computer Group, for their time and effort in assisting me during my project. Multiple times have I had your help with the installation of new software, and you gave me valuable advice on the use of the computing facilities, which really helped me to convert my raw data into results.

Furthermore, I would like to thank two of my friend groups, the "Boslovers" and the "Spain group", for keeping contact during the COVID lockdown periods, the support and interest you showed in my project, and the much needed social interaction via the various platforms we use, or in person when the lockdown measures allowed for that.

Needless to say, I would like to end by expressing my immeasurable gratitude to my parents and brother for their unconditional support throughout my school and university career. You have always showed enthusiasm for the things that have my interest and passion. This really motivated me to strive for the highest reachable, and it will hopefully continue to help me for the rest of my career.

Contents

1 Introduction

1.1 Summary

This master's thesis gives a full report of the findings from my research project, using radio interferometric observations of the radio galaxy Cygnus A with the high-band array of the Low-Frequency Array (LOFAR). Cygnus A is a radio galaxy of the brightest class, which form lobes of tens to hundreds of kiloparsecs in diameter that emit synchrotron radiation. This radiation, which is observable at radio wavelengths, is the result of the constant acceleration of the relativistic electrons in the lobes by the magnetic field of the surrounding cluster.

The plasma in the lobes is currently understood to find its origin in the two relativistic jets launched from the active black hole sitting in the central host galaxy. The content of these jet is diffused out and shock accelerated once it collides with the surrounding cluster gas. The location of these collisions can often be seen in the most powerful radio galaxies, including Cygnus A, in the form of bright circular regions near the edge of the lobe. These regions are named 'hotspots', and are still a big topic of research in astronomy as no conclusive model of their working has been found yet.

McKean et al. (2016) used a LOFAR observation between 110 and 180 MHz to directly observe a turnover in the spectra of the secondary hotspots of Cygnus A, which was the first observation of this phenomenon in Cygnus A. This turnover is expected to occur in regions such as hotspots, as these regions have a higher electron density and brightness temperature than the rest of the lobe, high enough to allow for partial (self-)absorption of the synchrotron emission produce in the lobe. However, a calculation of the physical conditions in the hotspots from the characteristics of the turnover did not yield parameter values that are in agreement with research on other observational evidence. However, the observations from the Dutch stations were of such good quality, that extending the observation to other LOFAR bands looked promising.

The observation used in this thesis consists of a new LOFAR data set on Cygnus A. The other frequency band of the high-band array was also used, adding the frequency range between 210 and 250 MHz. The calibration and imaging of the interferometric data resulted in an improved the angular resolution, allowing for a more detailed study of the structure of Cygnus A's radio lobes between 110 and 250 MHz. The intensity image reveals three of the four known hotspots, and more extended structure throughout the lobes that plausibly the result of shockwave producing events. Additional features can be seen in the spectral maps: The fourth hotspot is revealed, as well as a new region with hotspot-like characteristics. The spectral maps also show that the bright regions in the lobes have relatively flat spectra, suggesting that these regions are shock accelerated.

Spectra were taken from seven regions in Cygnus A, including hotspots A and D that were studied by McKean et al. (2016). These spectra were obtained from the LOFAR observation between 110 and 250 MHz, and three additional VLA images at 327 MHz, 1.037 GHz and 1424 GHz. Added to the list of regions are hotspot B (primary hotspot in the lobe), hotspot C (faint hotspot candidate in the counter-lobe) and three bright extended regions in the counter-lobe. Hotspot A and D show the previously observed turnover, with a signifcant part of the spectrum having a positive spectral slope. Hotspots B and C show a near turnover near 110 MHz, as their spectra are almost completely flat near the low-frequency end. Flattening in spectra of the extended regions in the counter-lobe is not as extreme as in the hotspots, but they show a significant amount of curvature, and combined with their relative brightness this suggests that these regions are shock

accelerated.

Modelling of synchrotron self-absorption and free-free absorption models to the hotspot spectra resulted in the same discrepancy experienced by McKean et al. (2016), as the magnetic field strength and electron density needed to cause the observed absorption are orders of magnitude higher than what is reasonable for such regions. A plausible solution has been proposed by Carilli et al. (1999): the hotspots are not homogeneous circular sources, but instead consist of smaller sub-structures that are not resolved by the LOFAR observation. This can potentially solve the magnetic field strength discrepancy, but observations with the international LOFAR stations are required to observe these sub-structures at LOFAR operating frequencies.

The curvature in the spectra of the extended regions is found to be unlikely due to aging of the electron populations that emit the synchrotron radiation, as the spectral age calculated from the curvature is unreasonably high. Shock acceleration could possibly be a better explanation of the curvature, but observations in other bands are needed to confirm this. International LOFAR observations will also be beneficial in understanding the origin of these regions, as observations of better angular resolution will allow for a more detailed look into their configuration and thereby contribute to their interpretation.

1.2 Active Galactic Nuclei

In this section the relevant background of Cygnus A, and AGN in general is given. It summarises the history of AGN, how they were discovered and how their interpretation developed, followed by an overview of their working using the unified model. Subsequently, the focus is pointed more towards radio galaxies, after which some recent findings of Cygnus A are discussed and their influence on AGN research.

1.2.1 Discovery of AGN

AGN are the brightest objects in the Universe currently known, sometimes outshining the stellar light of their host galaxy by several orders of magnitude. Their existence was only discovered in the second half of the 20th century, partially due to the lack of instrumentation suitable for observing AGN emission. Developments in astronomical imaging techniques in the early 20th century already opened up the possibility to identify AGN emission, but at the time the existence of galaxies outside the Milky Way was not even known, let alone the existence of AGN. Fath (1909) was the first to observe AGN emission in the form of emission lines in the spectra of Messier 77 and 81, which were then understood to be spiral nebulae inside the Milky Way, and Curtis (1918) observed a bright jet coming form Messier 87, which is shown in Figure 1. More mysterious emission spectra were observed in the following decades, but the realization that the emission was coming from other galaxies was only triggered after the discovery of galaxies outside the Milky Way by Hubble (1929). Several more AGN spectra were found in the following decades, most notably by Seyfert (1943), who found broad emission lines coming from galactic nuclei in a sample of nearby galaxies similar to Messier 77. This type of active galaxy, where the optical component of the host galaxy can be clearly detected, later obtained the name ”Seyfert galaxy”.

Breakthroughs in both radio and X-ray astronomy provided observations with better resolution at those wavelengths, making it possible to observe more distant sources. In particular, the Third Cambridge Catalogue of Radio Sources (abbreviated 3C, see Edge et al. (1959)) found many

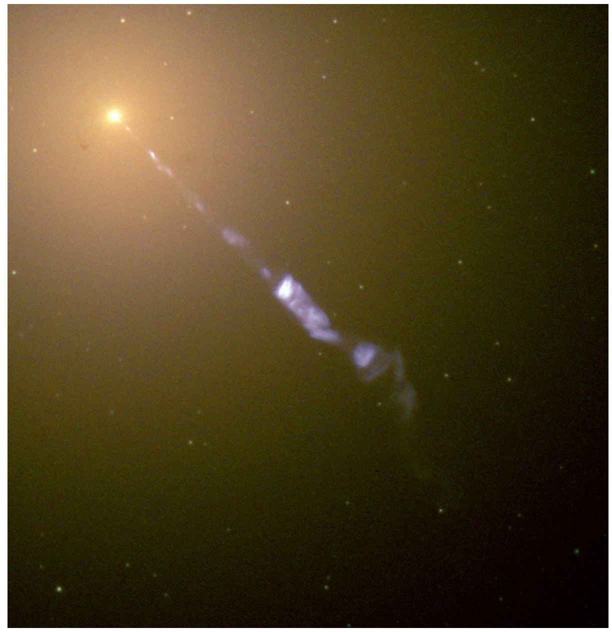

Figure 1: Composite image of M87 from the Hubble Space Telescope, showing a bright jet of plasma emerging from the center of M87. **Image credit: NASA and The Hubble Heritage Team (STScI/AURA)**

radio sources and made it possible to identify their optical components. Distant radio sources were found to be associated with pointlike optical sources, giving rise to the name of the other main AGN type: quasar (quasi-stellar radio source). The first real identification of a quasar was done by Schmidt (1963) when it was calculated that 3C 273 is so bright, and yet located so far from the Milky Way, that its emission can not possibly be powered by nuclear fusion.

A theoretical explanation was obtained independently by Salpeter (1964) and Zel'dovich (1964) a year later: The enormous power output can be explained as radiation resulting from accretion of gas by a supermassive black hole at the center of the galaxy. This process can convert the thermal and kinetic energy of the gas very efficiently, leading to enormous energy outputs. The theory has developed since then, adding more components as different examples of AGN show different spatial and spectral characteristics. The result is a still developing model stating that all AGN can be described by the same collection of components, and what is observed depends on their orientation. This 'unified model' is discussed in the next section.

1.2.2 AGN unification, components and types

Next to Seyfert galaxies and quasars, another handful of AGN types can be distinguished, each having their own spectral properties and appearance. It is theorized that all these types consist of the same components, but are observed differently due to their orientation with respect to the observer (Barthel, 1989; Antonucci, 1993). In this orientation-based unified model, an AGN consists of: a supermassive black hole, an accretion disk, a dense dusty torus, two high collimated relativistic jets of plasma, and a group of thin and dense clumps of gas (the "narrow line region" (NLR) and the "broad line region" (BLR), respectively). A schematic drawing of the arrangement of these components can be seen in Figure 2. This is the currently accepted model to describe all forms of AGN activity that are observed, as the varying behaviour of different AGN can be explained from the orientation towards the observer. The model fits most AGN observations, one recent and very important one being the direct observation of the accretion disk and photon sphere

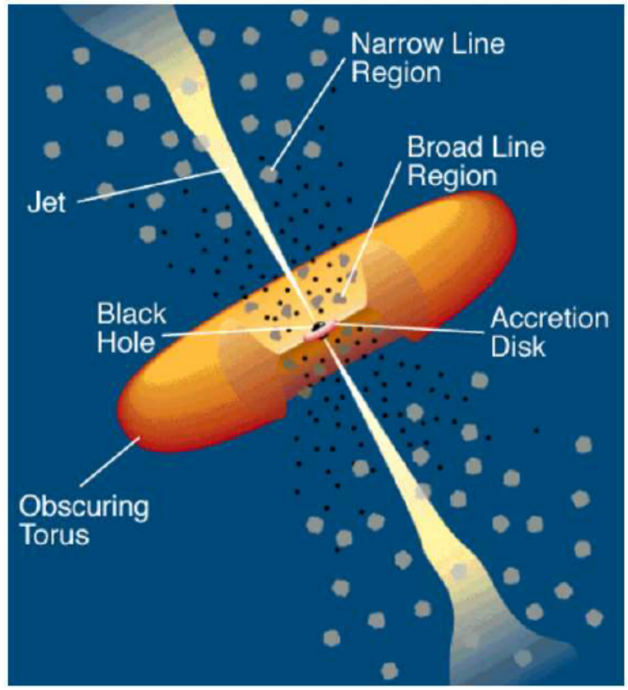

Figure 2: Schematic drawing (not to scale) from Urry and Padovani (1995) of the configuration of the currently favoured unified model of AGN. The black hole is surrounded by a hot radiating accretion disc, and a torus that obscures the central region if the AGN is not observed face-on. The jet is launched along the magnetic field lines around the AGN, perpendicular to the accretion disk plane. The BLR and NLR are found in and near the path of the jet.

of the central black hole in M87 (Akiyama et al., 2019).

The black hole is the central component of an AGN. Its mass can range from the order of $10^6 M_\odot$ (supermassive) up to $10^{10} M_\odot$ (ultramassive), with the heaviest black hole known today being the quasar TON 618, having a mass of $6.6 \times 10^{10} M_\odot$ (Shemmer et al., 2004). Around the black hole we find the accretion disk; a disk of orbiting gas and dust, shaped by the strong gravitation of the black hole. Accretion disks typically have a diameter in the order of 10^{-3} pc. Viscous friction between layers in the accretion disk cause the infall of matter onto the black hole, as it loses its angular momentum. This infalling matter adds to the mass of the black hole, increasing its size over time. At the same time this friction causes the material in the accretion disk to heat up, producing a lot of X-ray radiation by upscattering photons through the inverse Compton effect, further explained in Section 1.3.3.

The broad line region, as the name suggests, is the component responsible for the broad emission and absorption lines that can be seen in the UV and optical AGN spectra. The BLR consists of high density gas clumps that absorb X-ray and UV photons, which get re-emitted at UV and optical wavelengths. The clouds are located within ~ 1 pc from the black hole and have velocities up to 10^3 km/s, which causes the emission and absorption features to be heavily Doppler broadened.

The narrow line region is located much farther from the black hole than the broad line region, sitting at distances between $\sim$ 100-1000 pc. The NLR consists of very low density clouds, making that they allow for forbidden emission features, which are strongly represented in NLR spectra. Due to the low velocity of the clouds, these features experience little to no broadening and are thus sharply peaked.

A dense dusty torus can be found around the black hole, a donut shaped structure with inner and outer radii of the order of 1 to 100 pc, respectively. The torus is both physically and optically thick to radiation of all wavelengths, meaning that the radiation from the black hole region is obscured if the torus is in the line of sight between the black hole and the observer. The physical conditions in the torus can vary widely due to the amount of heating received from the black hole at different radii. The inner regions are warmed directly by the black hole and show signatures of ionization, while the outer regions are cold and sometimes allow for the existence of dust.

The brightest AGN often show one or two collimated relativistic jets of plasma, launched from the black hole along its magnetic field lines. The physical working behind the launching of the jet is still a topic of research, as the region from which the jet is launched is relatively small and therefore difficult to observe directly. A theory has been suggested by Blandford and Znajek (1977), in which energy and angular momentum are extracted from the rotation of the black hole. The orientation of the jets with respect to the observer determines whether both of them can be observed, as Doppler beaming will increase the luminosity of an approach jet, while a receding jet is less luminous through the same process (Blandford and Königl, 1979). The jets can have sizes of hundreds of kiloparsecs, reaching far out of the host galaxy. Jets can be observed from radio wavelengths in the form of synchrotron radiation, to X-rays because of the upscattering of photons through the inverse Compton process, similar to the accretion disk. Both these processes are discussed more in depth in Section 1.3.

1.2.3 Radio galaxies

Not all AGN are observed to form a jet: only radio-loud quasars, blazars and radio galaxies are observed to produce jets, the latter being the type that Cygnus A belongs to. The jet and its impact on the environment produces a lot of radio emission, making that AGN with jets are often classified as radio-loud. The radio-loudness of an AGN is defined as

$$R = \frac{L_{5\,\mathrm{GHz}}}{L_{\mathrm{B}}}, \tag{1}$$

where $L_{5\,\mathrm{GHz}}$ is the 5 GHz (radio) luminosity and L_{B} is the B-band (optical) luminosity. It is commonly accepted that the radio-loud AGN have $R > 10$, in contrast to radio-quiet AGN that usually have $R < 1$ (Visnovsky et al., 1992).

Blazars are observed face-on, i.e. their jet is pointed more or less towards the observer. Therefore their parsec-scale radio emission is often unresolved. For radio galaxies (and most quasars) this is not the case. Their jets are oriented at an angle, and hence they are projected on the sky on resolvable angular scales. This does not mean that radio galaxies always show two jets: the jet that is pointed slightly more towards the observer experiences Doppler beaming, increasing its luminosity significantly, while the jet pointing away will be dimmer through the same effect. This, together with the unbeamed jet brightness determines how many jets will be observed in a radio galaxy.

Next to the AGN and the jet, observations of radio galaxies show one more key feature: radio lobes. These are diffuse structures of scales between tens and hundreds of kiloparsecs that consist of synchrotron radiation emitting plasma. The structure of the radio lobes is what makes every radio galaxy unique, as they all have different shapes and sizes due to the different conditions in the intergalactic medium. However, they can still be divided further into two species by linking their overall morphology to their luminosity. This strategy was first proposed by Fanaroff and Riley (1974), and is known as the Fanaroff-Riley classification.

A few clear distinctions can be made between the morphology of less luminous (FR-I) and more luminous (FR-II) radio galaxies. The jet of an FR-I galaxy is at its brightest close to the AGN,

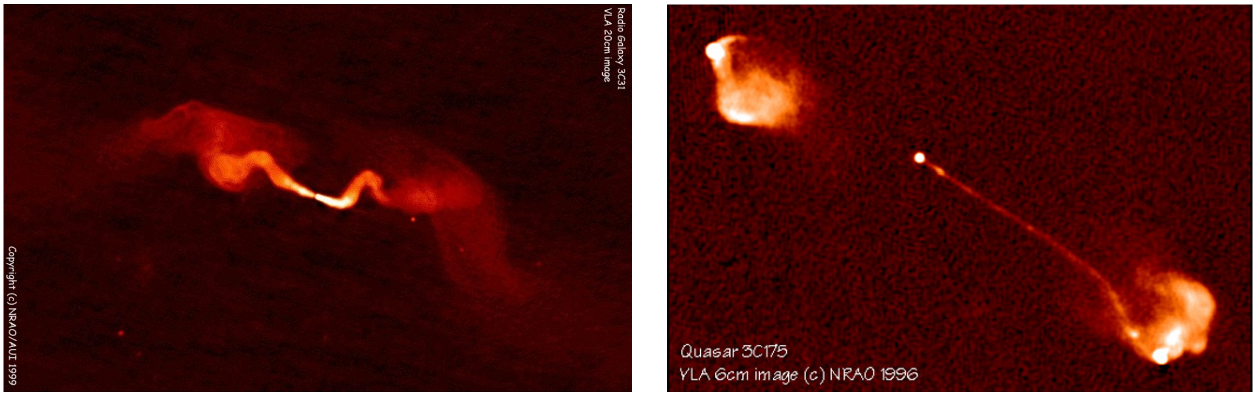

Figure 3: Left: VLA image of 3C31, a typical FR-I galaxy, at 20 cm. Two jets can be observed and are at their brightest close to the center. The two lobes are less bright, plume-like and start swerving further outward from the host galaxy. Right: VLA image of 3C175, a typical FR-II galaxy, at 6 cm. Only one of the jets can be seen due to relativistic beaming, and it is narrower than that of 3C31. Also two spherical lobes can be seen, being brightest at the edges in the hotspots and showing a lot of diffuse structure. **Image credit: NRAO**

and becomes less luminous further out from the center. The radio lobes are less bright than the jet, and often have a swerving, plume-like shape in the extension of the jet. It is for this reason that FR-I galaxies are also called "edge-darkened". Contrary, FR-II galaxies have much more luminous lobes, which are often fat and nearly spherical. Hence, FR-II galaxies are also called "edge-brightened". The lobes often contain "hotspots", which are bright compact structures where the jet comes to a halt as it collides with the surrounding cluster medium. These hotspots are still a topic of research, but they are currently understood to diffuse out the content of the jet through turbulence and shockwaves that move upstream through the jet. The diffused particles flow back from the hotspots to the AGN while getting accelerated by the magnetic field in the cluster, creating the diffuse synchrotron emitting lobes that keep expanding as energy keeps getting added into the lobes by the jet (Blandford and Rees, 1974).

The jet is a relatively weak emitter in FR-II galaxies and is sometimes outshone by the radio lobes completely, depending on the observing frequency. The nucleus of an FR-II galaxy can often be distinguished from the jets more easily than in FR-I galaxies, as it often radiates brighter than the relatively dim jets. An example of both radio galaxy types can be seen in Figure 3.

1.2.4 AGN feedback & outflows

The jets and radio lobes produced by AGN play an important role in galaxy formation theory, because they impact their host galaxy and its surroudings quite severely. Understanding these phenomena has shown to be of great importance, as it potentially solves one of the biggest problems in this field: the luminosity function.

The luminosity function is a distribution that describes the number density of galaxies in the universe per luminosity bin. The exact shape of this function depends on the behaviour of dark matter, the scaling relation between galaxy mass (inferred from luminosity) and dark matter halo

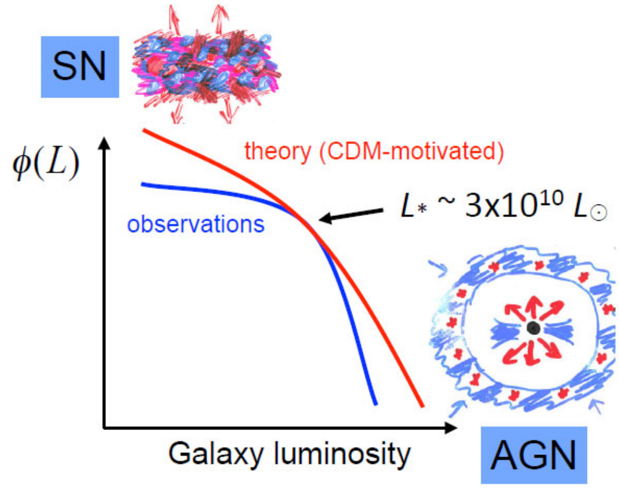

Figure 4: Schematic drawing of the discrepancy between theorized and observed luminosity function by Silk and Mamon (2012). The observed number density of galaxies is lower than expected for both the lowest and highest luminosity galaxies. The solutions are supernova feedback and AGN feedback, respectively.

mass, and the cooling behaviour of gas in clusters. However, after making educated assumptions on each of these, the resulting theoretical luminosity function still shows a discrepancy with observations: the theory heavily overestimates the number of galaxies of the lowest luminosities, as well as the galaxies with the highest luminosities, as visualized in Figure 4. The proposed mechanism to explain this discrepancy is the quenching of star formation before all gas in a galaxy has had the time to form stars. The source of this mechanism differs between low and high luminosity galaxies.

For low luminosity galaxies, the quenching is done by supernova feedback. Supernovae violently blow out the outer layers of a star at the end of its nuclear fusion lifetime. A supernova has two effects on the surrounding interstellar medium: it heats the gas, suppressing its cooling and thereby inhibiting star formation, and the expanding hot gas sweeps out more material in its path via ram pressure, disrupting and pushing out clumps of gas that are being kept from forming new stars. The impact of supernovae is so significant in small galaxies that it can disrupt the star formation process, while this effect is not as proportional in larger galaxies. The result is that the low luminosity galaxies are left as dim remnants that fall in a luminosity bin below the detection limit of the best telescopes, hence the low luminosity galaxy count is lower.

AGN feedback is the process responsible for the quenching of star formation in galaxies at the high luminosity end. The principle is similar to that of supernova feedback: the AGN can heat the gas in its surroundings, thereby stopping its cooling, and it can blow out material, stirring and relocating the surrounding gas clumps. This leaves the galaxy with less stellar mass than what was initially expected from its gas content.

The inhibiting of the cooling can be seen very clearly in X-ray images of clusters, and also in Cygnus A, in the form of cavities. X-ray images can be used to trace hot gas in clusters, which typically has temperatures up to 10^{10} K. Once this gas gets the chance to cool, it can clump together and form stars. X-ray images show that about half of all clusters have large cavities in the distribution of their gas content, which often perfectly coincide with the location of the radio lobes produced by the AGN in the cluster. This suggests that the radio lobes, inflated by the jet, have pushed out and heated the gas. As a result, this gas has not fallen into the galaxy and is kept from cooling to star forming temperatures, leaving a lot of loose gas in the cluster. An example of such a cluster, where the radio lobes have pushed out gas to form large X-ray cavities, can be seen in Figure 5. For a more extensive review of cluster gas heating and X-ray cavities, we refer

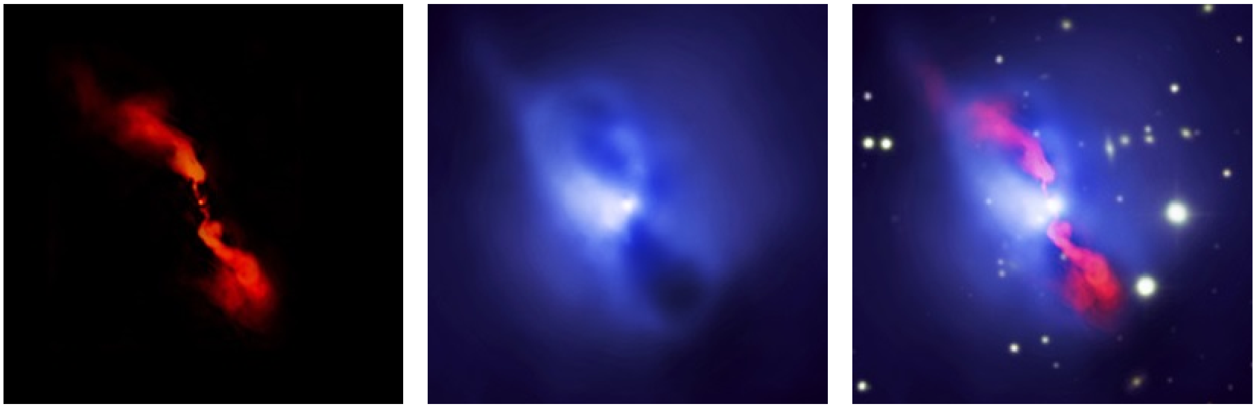

Figure 5: Radio (left), X-ray (center) and composite (right) image of galaxy cluster Hydra A. The composite image, consisting of radio, visible and X-ray continuum, clearly shows how the radio lobe plasma perfectly fits in the X-ray emission cavities, one of the key signatures of AGN feedback. **Images adapted from: X-ray: NASA/CXC/U.Waterloo/C.Kirkpatrick et al.; Radio: NSF/NRAO/VLA; Optical: CanadaFrance-Hawaii-Telescope/DSS.**

to McNamara and Nulsen (2007, 2012).

The blowout process can be seen on much smaller scales, in the form of outflows. Where the radio lobes and X-ray cavities have sizes of tens to hundreds of kiloparsecs, outflows are seen between the parsec and kiloparsec scale. Their existence and properties can be observed from absorption and emission lines, that get heavily blueshifted as the outflows can reach near-relativistic velocities. The outflow process is often very complex, and the phase of the outflowing gas can vary, depending on the physical circumstances. It is for this reason that outflows can be traced in emission and absorption lines of ionized gas, as well as neutral and molecular gas, spanning a wide observing frequency range.

Photons and jets from the AGN can cause outflows through two phenomena: the jet can push out material by exerting ram pressure on the gas in its path, or photons from the accretion disk can heat the surrounding gas via Compton scattering (thermal winds). The effect of the jet is rather simple: plasma from the jet is launched outward at high velocities, dragging along all the gas in its way. The heating is slightly more complicated: Compton heating by photons from the accretion disk causes hot bubbles of gas that shock the IGM as they expand through it, sweeping out the gas and dust in their path. Observations show that the outflowing material in both jet-driven and wind-driven outflows can contain neutral-atomic and molecular gas. This suggests that either the physical conditions allow the gas to cool radiatively, or that pre-existing molecular gas gets accelerated without getting destroyed.

The heating and outflowing of the gas caused by the AGN creates a negative feedback loop: when the AGN is fueled with material to accrete, its activity blows out the rest of the gas, getting rid of its own fuel. This causes the AGN activity to drop, and as a result the relocated and outblown gas can fall back to the AGN, fueling a new period of AGN activity. The active period of one iteration of this loop, a ”duty cycle”, lasts shorter, but is also more violent than the quiet period. The more massive galaxies usually host the AGN with the longer and more powerful duty cycles, explaining the shape of the high mass end of the luminosity function. Insights in the working of

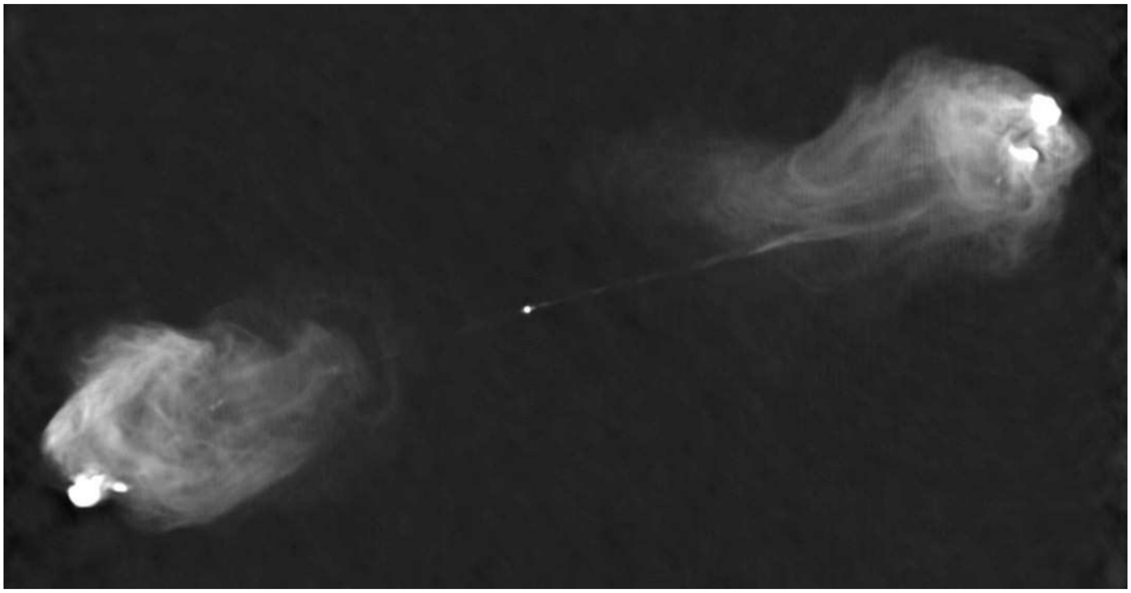

Figure 6: VLA radio image of Cygnus A at 5 GHz with 0.4" resolution from Perley et al. (1984). The dot in the center is the active black hole, and two relativistic jets are seen emerging in opposite directions. Both radio lobes show at least two hotspots in the extension of the jet, and the plasma flowing back from the hotspots shows to have formed a spherical lobe with a complex internal structure.

duty cycles are of great importance for our understanding of AGN fueling and jet formation, and thereby the resulting AGN feedback and its impact on galaxy formation. King and Pounds (2015) gives a more in-depth read into the working of outflows and the resulting AGN feedback process.

1.2.5 Cygnus A

Cygnus A is the textbook example of an FR-II radio galaxy. Being one of the brightest radio sources in the sky, its existence was discovered already in 1939 by radio pioneer Grote Reber. It appeared as one of the brightest sources in his radio map of the Milky Way, but the resolution was too poor to draw any conclusions on the source type. Using improved interferometric techniques in the following decades, it was discovered that Cygnus A is a double radio source (Jennison and Gupta, 1953). Together with the enormous energy output, this led to the hypothesis that it was a collision of two galaxies, as described by Baade and Minkowski (1954). The lack of a double optical source kept the debate on the origin of Cygnus A's radio emission open. A decade later, when the concept of AGN was postulated, researchers started to realise that Cygnus A is an ultra-luminous radio galaxy. This was supported by a calculation by Hargrave and Ryle (1974), who concluded that jet activity was necessary to explain the existence of the hotspots in the radio lobes. The first observation of the jet and thereby the confirmation of the jet powered model was done by Perley et al. (1984) using the Karl G. Jansky Very Large Array (VLA). A radio image from this observation can be seen in Figure 6.

Cygnus A is by far the most luminous 'local' radio galaxy. Most radio galaxies of comparable luminosity are found at significantly larger distances, around redshift $z \sim 1$, whereas Cygnus A is located at $z \sim 0.056$. This makes Cygnus A the perfect facility to study FR-II galaxies in detail and get an insight in their working, as its size on the sky is much larger than that of comparable galaxies. Observations across the electromagnetic spectrum show many spectral and spatial features in the jet, lobes and hotspots that feed ongoing research into active galaxies and help understanding how they fit in the astrophysical framework.

One of the clearest characteristics that can seen in observations of Cygnus A, such as Figure 6, is that Cygnus A contains multiple hotspots in both lobes. The most widely accepted theory is that the primary hotspots are currently being fed with plasma by the jet, while other hotspots are relic hotspots, which cool by expansion. This explanation, proposed by Scheuer (1982), is known as the "dentist drill model". The visualization of this model is that the jet can change direction over time due to precession or a jet deflecting event, causing it to hit the inner edge of the X-ray cavity at different locations over time. One recent study by Steenbrugge and Blundell (2008) backs up this theory, as it is found that the jet can be seen swerving into the lobe, possibly due to precession, and the counter-jet is seen to deviate from a straight line when entering the counter-lobe. This behaviour requires the speed of the jet to be relatively low, which is verified by multiple independent calculations. It is also consistent with the finding from Pyrzas et al. (2015) that the secondary hotspot in the lobe has an age of the order of the precession period.

In the same series of publications it is also found that X-ray images of the counter-lobe show characteristics of a relic jet and relic lobe plasma, pointing towards activity of a previous duty cycle (Steenbrugge et al., 2008, 2010). Other evidence for earlier activity was found by Chon et al. (2012), who showed that an X-ray cavity filled with synchrotron emitting plasma can be found towards the south of the nucleus, which is likely the result of an earlier duty cycle with a jet axis pointing in a completely different direction than its current direction. This behaviour has been seen in more radio galaxies, showing radio structures of significant size along multiple jet axes. These galaxies have been given the name "winged" or "X-shaped" radio galaxies (Hogbom and Carlsson, 1974; Ekers et al., 1978; Leahy and Williams, 1984).

Furthermore, it was shown by Steenbrugge et al. (2014); Snios et al. (2018) that the nucleus of the AGN is offset and moving away from the center of the radio structure, likely due to proper motion of the host galaxy within the cluster. The nucleus also has a companion radio source, named Cygnus A-2, orbiting around the supermassive black hole. The nature of the source is uncertain, but its brightness makes it likely an extremely powerful supernova or another supermassive black hole (Perley et al., 2017).

1.3 Emission & absorption mechanisms

This section gives a physical background on the most important mechanisms that are responsible for producing the emission seen on the larger scale in Cygnus A. The most dominant mechanism in the radio lobes is synchrotron radiation, combined with inverse Compton scattering. Synchrotron emission is seen brightly at radio wavelengths, while inverse Compton scattered radiation is represented by X-rays. A series of formulae and proportionalities will be derived, although some parts of the derivation will not be done explicitly. For a more detailed explanation and the complete derivation we refer to Rybicki and Lightman (2008).

1.3.1 Synchrotron emission

The radio emission in the lobes of radio galaxies is dominated by the synchrotron process. This is the relativistic version of the cyclotron process: particles in a magnetic field are accelerated by the Lorentz force, causing them to emit radiation. Because the Lorentz force acts perpendicular to both the direction of the magnetic field and the velocity vector of the particles, they follow a helical path around the magnetic field lines.

The gyrofrequency of a relativistic electron with mass m_e and charge e moving perpedicular to a magnetic field of strength B can be calculated from balancing the centrifugal and magnetic force acting on the electron:

$$m_e\omega^2 r = \frac{e\omega r B}{\gamma c} \implies \omega_{\mathrm{B}} = \frac{eB}{\gamma m_e c} \implies \omega_{\mathrm{G}} = \frac{eB}{m_e c}, \tag{2}$$

where r is the gyration orbit radius and $\gamma = 1/\sqrt{1 - v^2/c^2}$ is the Lorentz factor of the electron. Using numerical values, this relation becomes

$$\frac{\omega_{\mathrm{G}}}{\mathrm{MHz}} = 17.6\left(\frac{B}{\mathrm{Gauss}}\right). \tag{3}$$

This is an angular frequency, hence an electron completes one orbit in

$$\frac{\nu_{\mathrm{G}}}{\mathrm{MHz}} = 2.8\left(\frac{B}{\mathrm{Gauss}}\right), \tag{4}$$

where ν_{G} is the gyration frequency in orbits per second. The synchrotron power radiated by a relativistic electron in such an orbit is given by the Larmor formula

$$P' = \frac{2e}{3c^3}\left(a'_\perp\right)^2, \tag{5}$$

where the tangential acceleration $a'_\perp$ and power P' are measured in the inertial frame of the electron. Using Lorentz transformations it can be shown that these quantities transform to the rest frame of the galaxy (which is approximately at rest for the observer) as

$$P = P' \quad \text{and} \quad a_\perp = \frac{a'_\perp}{\gamma^2}, \tag{6}$$

therefore the observed power is

$$P = \frac{2e\gamma^4}{3c^3}(a_\perp)^2. \tag{7}$$

Using the same force balance applied in Equation (2), the acceleration can be rewritten as

$$a_\perp = \frac{eBv_\perp}{\gamma m_e c} = \frac{eBv\sin\alpha}{\gamma m_e c}, \tag{8}$$

where $\alpha = \frac{v_\perp}{v_\parallel}$ is the 'pitch angle' of the electron. Averaging the power over the synchrotron lifetime of the electron, we obtain

$$\langle P_{\mathrm{sync}}\rangle = \frac{4}{3}\sigma_{\mathrm{T}}\beta^2\gamma^2 c U_B, \tag{9}$$

where $\beta = \frac{v}{c}$, $\sigma_\mathrm{T} = \frac{8\pi}{3}\left(\frac{e^2}{m_e^2 c}\right)^2$ is the Thomson cross section and $U_B = \frac{B^2}{8\pi}$ is the magnetic energy density. The synchrotron lifetime of the electron can then be expressed as

$$t_\mathrm{sync} = \frac{E_e}{\langle P_\mathrm{sync}\rangle} = \frac{\gamma m_e c^2}{\frac{4}{3}\sigma_\mathrm{T}\beta^2\gamma^2 c U_B} \approx 16.4\,\mathrm{yrs}\cdot\frac{1}{\gamma}\left(\frac{1\,\mathrm{Gauss}}{B}\right)^2, \tag{10}$$

where E_e is the electron's energy. This tells us that both higher Lorentz factors (higher kinetic energy) and stronger magnetic fields result in a shorter synchrotron lifetime of the electron.

An important point to note is that the synchrotron radiation from one single electron will not be seen continuously by an observer. This is because the relativistic speed of the electron causes its emitted radiation to be Doppler beamed into a narrow cone around its velocity vector. Hence, the observer will only see the radiation when they are within the 'opening angle' of the electron, which is approximately $1/\gamma$ in the relativistic limit. It is for this reason that the observer sees pulses of radiation coming from a single electron. The observed power spectrum from that single electron, which is more or less continuous, can be obtained by doing a Fourier transform of the pulse series, becoming

$$P(\nu) = \frac{\sqrt{3}e^3 B\sin\alpha}{m_e c^2}\left(\frac{\nu}{\nu_c}\right)\int_{\nu/\nu_c}^{\infty} K_{5/3}(\eta)\,\mathrm{d}\eta\,, \tag{11}$$

where $\nu_c = \frac{3}{2}\gamma^2\nu_G\sin\alpha$ is the critical frequency, where most of the electron's energy is radiated, $\int_{\nu/\nu_c}^{\infty} K_{5/3}(\eta)\,\mathrm{d}\eta$ is the synchrotron spectrum and $K_{5/3}$ is a modified Bessel function.

Knowing the exact shape of the synchrotron spectrum emitted by a single electron is not essential to find the shape of the synchrotron spectrum in astrophysical sources, as they contain large amounts of electrons, each having their own energy and pitch angle. Each of these electrons radiates most of its energy near the critical frequency ν_c, making that their contribution at other frequencies is outshone by electrons having those critical frequencies. The electron energy distribution of an electron ensemble in an astrophysical body generally follows a power law, given by

$$N(E)\,\mathrm{d}E = CE^{-\delta}\,\mathrm{d}E\,, \tag{12}$$

where $N(E)\,\mathrm{d}E$ is the volume density of electrons with energy between E and $E + \mathrm{d}E$, δ is the "injection index", and C is a normalization constant. The emission coefficient is then given by

$$\epsilon_\nu\,\mathrm{d}\nu = -\frac{\mathrm{d}E}{\mathrm{d}t}N(E)\,\mathrm{d}E\,, \tag{13}$$

which scales as

$$\epsilon_\nu \propto \nu^{(1-\delta)/2}. \tag{14}$$

The flux density S_ν is proportional to the emission coefficient, and therefore also scales as

$$S_\nu \propto \nu^\alpha, \tag{15}$$

where $\alpha = (1-\delta)/2$ is defined as the spectral index. The important conclusion to take away from this derivation is that an ensemble of electrons having a continuous energy distribution will produce a spectrum consisting of the sum of all individual electron spectra, to form a power law that does not depend on the shape of the spectrum of the individual electrons. A visualization is given in Figure 7.

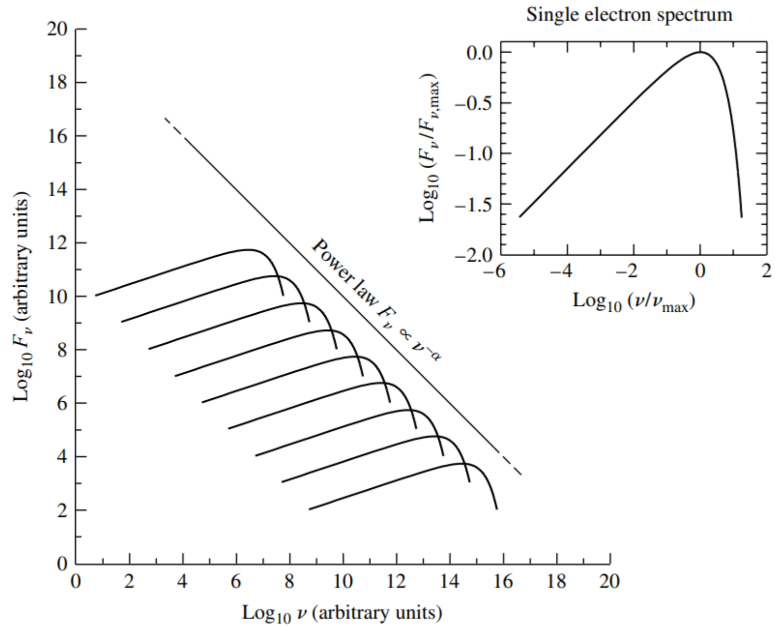

Figure 7: Visualization of synchrotron power spectrum from an ensemble of electrons from Carroll and Ostlie (2017), showing how the individual electron power spectra (top right figure) sum up to form a power law spectrum. Not shown are the turnover from absorption processes at low frequencies and the steepening due to aging at high frequencies. Note that the convention $F_\nu \propto \nu^{-\alpha}$ is used in this figure.

1.3.2 Components of a synchrotron spectrum

In the previous section it was shown that a continuous distribution of electron energies in an astrophysical source will produce a power law spectrum, where the shape depends on the electron energy distribution, and not the shape of the individual electron spectra. However, it has been theorized and observed in synchrotron sources, including Cygnus A, that the spectral behaviour varies depending on the observing frequency under the influences of absorption processes and radiation losses (Kardashev, 1962; Scheuer and Williams, 1968; Rybicki and Lightman, 2008; Longair, 2010). Spectra of synchrotron sources can often be divided into two characteristic frequency regimes that have opposite spectral slopes compared to one another.

The shape of the low frequency spectrum has been theorized, but only a few observations have been done so far of this regime in radio galaxies as sensitive interferometers with long baselines are needed to obtain a reasonable spatial resolution at these frequencies. The theory predicts that the low frequency synchrotron spectrum experiences synchrotron self-absorption (SSA). Synchrotron photons of these frequencies are more likely to be absorbed by another synchrotron electron within its source, rather than leaving the source unabsorbed. In other words, the source itself is optically thick to low frequency synchrotron photons, making that most of the observed radiation at these frequencies originates from electrons near the edge of the source. This behaviour can more or less be interpreted as a black body, although synchrotron emission itself is non-thermal. For this reason, the spectral behaviour shows to be similar, but not exactly equal to the Rayleigh-Jeans law ($B_\nu \propto \nu^2$) seen in thermal black bodies. The low frequency end of the synchrotron spectrum has a spectral index value of $\alpha = 5/2$, i.e. the flux density goes as $S_\nu \propto \nu^{5/2}$.

Other phenomena have been proposed to play a role at low frequencies as well, because fitting SSA models to synchrotron spectra does not always give a satisfactory result, as the optically thick spectrum often seems to be flatter than $\alpha = 5/2$. SSA fitting also requires physical parameters that are orders of magnitude off from the realistic values in the case in Cygnus A's hotspots (Harris et al., 1994; Carilli et al., 1999). For this reason, a different phenomenon responsible for creating the positive spectral slope in Cygnus A has been proposed by Carilli et al. (1991), being a low energy cut-off (LEC) in electron energies. This LEC is theorized to be the result of the conversion of the bulk kinetic energy of the jet into kinetic energy and turbulence in the electrons diffusing out from the hotspots, causing nearly all electrons to have more energy than a specific threshold (Godfrey et al., 2009). If an LEC is the dominant feature at low frequencies, the low frequency spectrum is dominated by the low frequency spectrum of the electrons with energies just above the cutoff. This single electron spectrum is found to have a spectral index of $1/3$, and therefore the LEC is expected to form a low frequency synchrotron spectrum of the form $S_\nu \propto \nu^{1/3}$. This slope is much flatter than that produced by SSA, and thus the distinction between the two can clearly be made if either SSA or an LEC would be the dominant feature. However, it will be shown later in this thesis that the spectra in Cygnus A will not be able to completely rule out either of the processes, hinting towards a third process that is at play.

Free-free absorption (FFA) has also been proposed as a suitable absorption mechanism that can effect the low frequency spectrum. FFA is the process in which energy from passing photons gets absorbed by an electron in an ionized medium along the line of sight between the observer and the source, or in the source itself. The low frequency spectrum created by FFA lies somewhere between the slopes resulting from LEC and SSA, and the full shape of the spectrum is determined from the optical depth of the free-free absorbing cloud (Kassim, 1989).

At higher synchrotron frequencies, the source is optically thin to the synchrotron photons. In this regime the shape of the synchrotron spectrum is heavily determined by the energy distribution of the synchrotron emitting electrons, as discussed in Section 1.3.1. Observations show that the optically thin synchrotron spectrum in most astrophysical sources has a slope $-1.25 < \alpha < 0$ (Kellermann, 1964). The exact value varies from source to source, and within sources as well, as is the case in Cygnus A (Kassim et al., 1996; Lazio et al., 2006). This is because the electron energy distribution can change throughout the source under the influence of two factors: the magnetic field strength and the cooling of the synchrotron electrons. These two together form the process known as spectral aging. As derived in Equation (10), high energy synchrotron electrons radiate away their energy at a faster rate than lower energy electrons, and therefore their synchrotron frequency falls off more rapidly compared to electrons in the same population. This has two consequences for the appearance of an aged synchrotron spectrum: a break arises at the high frequency end where the flux density starts to fall off exponentially, and the rest of the spectrum becomes steeper over time as all cooled electrons lose their kinetic energy, lowering their emitted synchrotron frequency. The break frequency is often used in spectral age modelling, as it is known to scale as

$$\left[\frac{\nu_{\mathrm{B}}}{\mathrm{GHz}}\right] = 1610^2 \left[\frac{B}{\mu\mathrm{G}}\right]^{-3} \left[\frac{\tau_{\mathrm{sync}}}{\mathrm{Myr}}\right]^{-2} \tag{16}$$

where B is the magnetic field strength, and τ_{sync} is the age of the electron population, i.e. the time that has passed since the spectrum was a power law stretching out to infinity (Carilli et al., 1991; Parma et al., 1998; Blundell and Rawlings, 2002). A schematic drawing of the aging of a spectrum can be seen in Figure 8. In FR-II radio galaxies, synchrotron emission close to the nucleus is expected to be radiated by the oldest electron populations, as a long time has past since these electrons diffused out from the hotspot. Hence, these electrons have experienced a lot of cooling and the break frequency in these regions is found at lower frequencies than close to the hotspots.

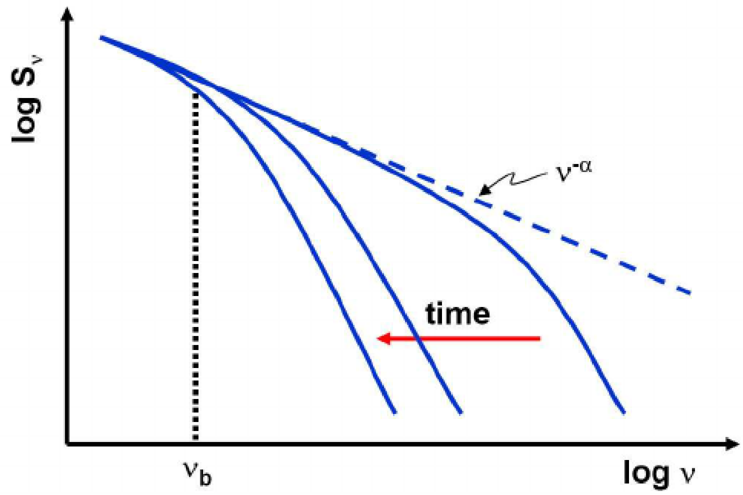

Figure 8: Schematic drawing of the aging of a synchrotron spectrum from Klein (2008). Break frequency ν_b is the point where the slope of the spectrum changes from a power law to an exponential drop-off due to the cutoff in high electron energies. The break shifts towards lower frequencies as the electron population ages.

Between the optically thick and optically thin regime, there is a point where the slope of the spectrum turns over from a positive spectral index to a negative index. This part of the spectrum is generally referred to as the ”spectral turnover”. Observing spectral turnovers directly is of great importance in the field of radio astronomy, because the shape of the full spectrum can constrain which of the previously mentioned absorption processes are at play in radio sources, and fitting models of these processes can give indications of the physical conditions in the sources. The turnover in the secondary hotspots of Cygnus A has been observed before by McKean et al. (2016), facing exactly the issue that the turnovers are too steep at low frequencies to be explained only by an LEC, while the fitted high frequency SSA power law becomes too flat. Both SSA and FFA models also result in unphysical parameters concerning the magnetic field strength in the hotspots and the electron density in the absorbing cloud, even if FFA happens in the hotspot itself.

1.3.3 Inverse Compton effect

Earlier we saw that synchrotron photons produced in radio galaxies like Cygnus A can be scattered by other electrons, the phenomenon known as synchrotron self-absorption. This process is a specific case of the more generally known ”Compton effect”, which is the inelastic scattering of a photon by a charged particle (mostly electrons) that converts part of the photon’s energy into the particle’s kinetic energy, hence lowering the energy of the photon. If the energy (either kinetic or thermal) of the particle is sufficiently high, the inverse process can also occur, i.e. the photon absorbs energy from the electron, upscattering it to a shorter wavelength. The ”inverse Compton effect” (IC) is known to occur in astrophysical sources that contain high amounts of relativistic electrons, such as Cygnus A. The result of this effect is that the synchrotron photons produced in Cygnus A will gain in energy up to X-ray energies by multiple interactions with other relativistic electrons, which is known as synchrotron self-Compton (SSC). This process also acts as an extra electron cooling process, as the electrons transfer part of their energy to every photon that they upscatter.

It can be shown, taking into account relativistic effects for the electrons, that the net power gain of the radiation field by the IC process is given by

$$P_{\mathrm{IC}} = \frac{4}{3}\sigma_{\mathrm{T}} c \beta^2 \gamma^2 U_{\mathrm{rad}}, \tag{17}$$

where U_{rad} is the initial radiation field, which can also be interpreted as the initial photon energy density. It should be noted that this equation is analogous to the calculation of the emitted synchrotron power P_{sync} in Equation (9), where the difference is that P_{IC} depends on the radiation field U_{rad}, and not the magnetic field density U_B. Combining Equation (9) and Equation (17) holds

$$\frac{P_{\mathrm{IC}}}{P_{\mathrm{sync}}} = \frac{U_{\mathrm{rad}}}{U_B}, \tag{18}$$

from which it can be concluded that the emitted IC power heavily depends on the relativistic electron density, as the electron density adds to the emitted synchrotron power but thereby also to the radiation field. From this we can conclude that SSC emission in Cygnus A will be observed in the regions with high electron and photon densities, i.e. the nucleus, the jets, the hotspots and possibly shock waves.

For an electron energy distribution of the form described by Equation (12), the SSC spectrum is found to depend on the frequency as

$$S_\nu \propto \nu^{(1-\delta)/2} = \nu^{\alpha}, \tag{19}$$

which is exactly the same dependence as for the optically thin synchrotron regime. An SSC spectrum therefore looks similar to an optically thin synchrotron, but at X-ray wavelengths.

2 Interferometry & Aperture Synthesis

This chapter gives a short review of interferometry, which is essential to understand the working of the LOFAR telescope and the description of the data reduction in Section 3. It starts off with the discussion of a simple two-element interferometer, followed by the generalisation to aperture arrays and the algorithms used in the processing of observational data to form images and spectra.

2.1 Simple interferometry

Astronomers have always searched for ways to obtain higher angular resolution in their images, and especially in the radio regime this is important, as the resolution of a single telescope is limited by the observing wavelength and the aperture size. The resolution of a telescope can be approximated by

$$\theta \approx \frac{\lambda}{D}, \tag{20}$$

where θ is the angular resolution in radians, and λ and D are the observing wavelength and the aperture diameter, respectively. It can be seen from Equation (20) that the resolution becomes very poor at long wavelengths.

The solution to this problem is measuring interferometric fringes using multiple telescopes as an interferometer, increasing the virtual size of the telescope from the diameter of a single antenna to the distance between the antennas. The resolution is then given by

$$\theta \approx \frac{\lambda}{B}, \tag{21}$$

where B is the distance between two antennas, which is called a "baseline". One of the first astronomical observations using an interferometer was done by Michelson and Pease (1921), who used a two element interferometer to measure the size of the red supergiant Betelgeuse. A schematic drawing of the setup and the result can be seen in Figure 9. The amplitude of the fringes is an indication of the resolving power of the interferometer, as any resolved portions of Betelgeuse cause destructive interference, dampening the fringes. They defined a quantity known as the interferometric visibility, given by

$$\mathcal{V} = \frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}}, \tag{22}$$

where $I_{\max}$ and $I_{\min}$ are the maximum and minimum of the peak fringe, respectively. From their calculation of $\mathcal{V}$ for different baselines they could calculate at what baseline Betelgeuse was resolved, thereby knowing its angular diameter.

2.2 Van Cittert-Zernike theorem & uv-coverage

The observation by Michelson and Pease (1921) measured only one visibility value, which they could use to calculate Betelgeuse's size only because of the assumption on the shape of their source. In reality, a source does not produce one visibility value, but instead gives a function of visibility values that depends on the observing baseline. For a distant incoherent monochromatic

Figure 9: (**a**) Schematic drawing of the Michelson-Pease stellar interferometer. Incoming light from a_1 and b_1 travels an equal distance to the detector, while light from a_2 arrives with a delay Δ w.r.t. light from b_2. The projected baseline changes by a factor of $\sin\theta$ when the source moves an angle of θ on the sky. (**b**) Schematic representation of the fringe pattern as an unresolved star ($\mathcal{V}$ = 1.0, solid curve) and a partially resolved star ($\mathcal{V}$ = 0.5, broken curve) move through the beam of the (fixed) apertures. (Thompson et al., 1986)

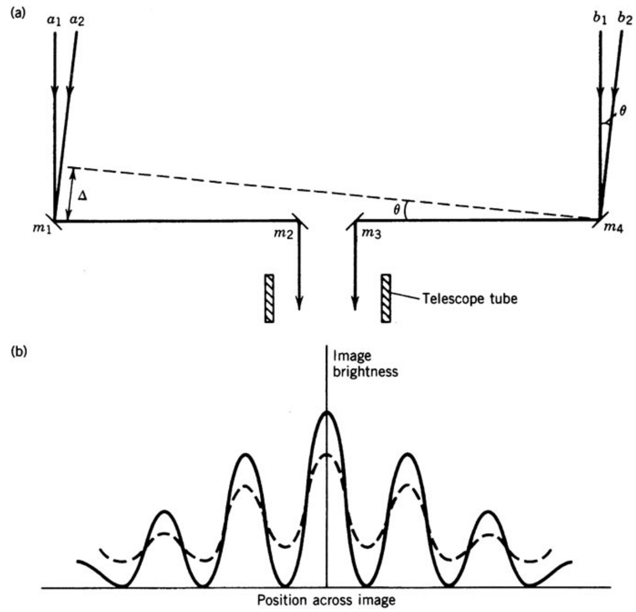

source, this visibility function can be described by a Fourier transform in the uv-plane, i.e. the coordinates of the interferometer, known as the Van-Cittert Zernike theorem (Zernike, 1938):

$$\mathcal{V}(u, v, w \approx 0) = \iint I_\nu(l, m) e^{-2\pi i(ul+vm)}\, \mathrm{d}l\, \mathrm{d}m\,. \tag{23}$$

Here $I_\nu(l, m)$ is the brightness distribution of the source at sky coordinates l and m, and u, v and w are the relative Earth coordinates of the interferometer, measured in number of observing wavelengths λ. We can invert the relation to become

$$I_\nu(l, m) = \iint \mathcal{V}(u, v) e^{+2\pi i(ul+vm)}\, \mathrm{d}u\, \mathrm{d}v\,, \tag{24}$$

which tells us that we can retrieve the luminosity distribution of our source by simply doing an inverse Fourier transform of the observed interferometric visibilties. That is, if the visibility values at all uv-coordinates are known.

In reality, the number of visibilities that can be measured (and thereby the number of spatial scales that can be resolved at the same instance) is limited by the number of baselines, which is in turn constrained by the number of used antennas. In an observation, all baselines trace out an arc in the uv-plane as the Earth's rotation around its axis changes the orientation of the baselines w.r.t. the source, measuring a visibility value at every uv-coordinate they pass. Due to the resulting incomplete uv-coverage, the resulting observation is not perfect, introducing the other component of the observed brightness distribution: the "dirty beam". The dirty beam is the response of the array to a point source, and is calculated by doing a Fourier transform of the uv-coverage of the observation:

$$B(l, m) = \iint S(u, v) e^{2\pi i(ul+vm)}\, \mathrm{d}u\, \mathrm{d}v\,, \tag{25}$$

where $S(u, v)$ is the sampling function in the uv-plane of the interferometer. The observed brightness distribution can then be written as a convolution of the real brightness distribution and the dirty beam of the array:

$$I_{\text{measured}}(l,m) = I_{\text{real}}(l,m) \otimes B(l,m). \tag{26}$$

This is also known as the "dirty image", which is the response of an interferometer to a complete source. For extended sources the dirty image can become quite messy, as every piece of real structure is smeared over the whole image by the point spread function of the array. The CLEAN deconvolution algorithm is used to retrieve an approximation of the real brightness distribution of the source. This is described in detail in Section 2.4. However, the observed data must be calibrated before imaging can be done.

2.3 Self-calibration

Uncalibrated observations contain errors in the visibility measurements, due to changing antenna gains during the observation. This is caused by two effects: instrumental and atmospheric. Examples of instrumental effects are variations in time of the antenna temperatures and instrumental delays. Atmospheric errors arise due to differences in the temperature and electron content of the atmosphere at different antenna locations (mostly on long baselines). These errors show up in what is known as the Hamaker-Bregman-Sault measurement equation, given by

$$\vec{\mathcal{V}}_{ij}^{\text{meas}} = J_{ij}\vec{\mathcal{V}}_{ij}^{\text{ideal}}, \tag{27}$$

where $\vec{\mathcal{V}}_{ij}^{\text{meas}}$ and $\vec{\mathcal{V}}_{ij}^{\text{ideal}}$ are the measured and true visibility of baseline ij at a certain sample time and spectral channel, and J_{ij} is the Jones matrix containing the intstrumental and atmospheric errors (Hamaker et al., 1996). Here the visibilities are a vector containing 4 polarizations, and hence J_{ij} is a 4×4 matrix. All errors embedded in the Jones matrix can be corrected for by calibrating the relative antenna gains over the observation, which will increase the quality and accuracy of the observation.

Calibration can be done during the observation itself, by switching to other sources of known surface brightness during the observation to monitor the atmospheric behaviour and the instrumental variations between measuring intervals. After these corrections, the observation can still be polished further by the process of self-calibration. In this technique, images of the observation itself are used to improve the image quality. When a model image is obtained from the observational data (described in detail in Section 2.4), the corresponding visibility values at each uv-coordinate can be calculated by a Fourier transform. The antenna gains during the observation can then be calibrated so that the observed visibilities best match the model visibilities, by a least-squares method. This process minimizes the difference between the observed visibilities $\mathcal{V}_{ij}^{\text{meas}}$ and the model visibilities $\mathcal{V}_{ij}^{\text{model}}$, given by

$$\sum_{\text{time}} \sum_{i<j} w_{ij} \left| \mathcal{V}_{ij}^{\text{meas}} - g_i g_j^* \mathcal{V}_{ij}^{\text{model}} \right|, \tag{28}$$

where g_i and g_j^* are the respective gains of antenna i and j, and w_{ij} is a weighting factor to favour visibilities with low variances (Thompson et al., 1986). This may look a bit dubious, as it seems as if an incorrect model can force an incorrect calibration on the visibility set, destroying the observation. While this is possible, this can often be done safely for interferometers with many antennas because the system is heavily overdetermined: every observational interval uses N antennas and thus needs to solve for N antenna gains, while $\frac{N(N-1)}{2}$ visibility values are measured, namely one for every baseline. Having $\frac{N(N-1)}{2}$ measurements and N unknowns allows to find many redundant solutions that should all be in line with each other. That is, if the initial observation is of reasonable accuracy. Another view of the situation is that a reasonable combination of antenna gains is found to produce the observed visibilities, instead of forcing the visibilities to their desired values. Usually multiple rounds of self-calibration and imaging are done, until there is no significant increase in image quality anymore. An example of gain corrections produced by self-calibration can be seen later in Figure 17.

2.4 CLEAN algorithm

After the calibration of the antenna gains in an attempt to obtain the ideal visibilities, an inverse Fourier transform of the observed visibilities can be done to obtain the dirty image. To recover the real brightness distribution from the dirty image, the adjacent lobes of the point spread function need to be scraped off the image, while keeping the real components intact. Since both the dirty beam and the observed visibilities are Fourier transforms, we can make use of the convolution theorem, which states that the Fourier transform of a convolution of two functions is equal to the product of the respective Fourier transforms. For the interferometer, this becomes

$$\mathcal{F}(B(l,m) \otimes I(l,m)) = \mathcal{F}(B(l,m)) \cdot \mathcal{F}(I(l,m)). \tag{29}$$

From the convolution theorem, we deduce that the observed visibilities can be approached as a linear product of the dirty beam and the brightness distribution in Fourier space. The big advantage of this is that the sidelobe structure in the dirty image can be removed by simply taking out the sidelobe structure for each point source in the dirty image. This is done by CLEAN, introduced by Högbom (1974). The CLEAN algorithm is an iterative deconvolution process that approximates the observed brightness distribution from the dirty image by performing the following steps:

1. Do a Fourier transform of the observed visibilities to calculate the dirty image, which still contains all the sidelobes from the dirty beam.
2. Find the location of the peak intensity in the dirty image and subtract the dirty beam from the image at this location. The amount of flux that is subtracted is a fraction (the 'loop gain' γ, usually $\sim$ 0.1), times the peak intensity. Put a delta function with an amplitude of the subtracted dirty beam at the same coordinate in a model image.
3. Repeat step 2 until there is no real structure left in the residual dirty image. Determining the end of the iterative process can be done in various ways, including:
   - Setting an RMS noise value in the residual image at which iterations will stop once it is reached.
   - Setting a maximum amount of iterations.
   - Setting a certain ratio of the peak intensity and the RMS noise in the residual image at which iterating will stop.
   - Stopping iterations once an increase in RMS noise between two iterations happens.
4. Convolve each delta function in the model image with a 'clean beam', an approximation of the main lobe of the dirty beam, usually a Gaussian with the same full width half maximum (FWHM).

5. Add the residual image from step 3 to the image from step 4 to obtain the 'clean image'.

These steps often produce decent results in the cleaning of dirty images that contain one or multiple point sources. However, the convergence of the algorithm can be rather slow for extended sources like Cygnus A, as structures spanning multiple pixels are modelled pixel by pixel. In addition, instabilities in the model are known to arise for this reason as well (Cornwell, 1983). To obtain a better source model, the multi-scale CLEAN algorithm can be used (Cornwell, 2008). This process follows the steps described above, with the difference that the dirty image is searched for multiple structure sizes, instead of searching delta functions only. These structure sizes are specified by the user, and can be interpreted as truncated Gaussian functions of the specified scales. The net result is that step 2 is extended to support multiple angular scales; the dirty image is searched for the peak residual, taking all specified scales into account with their respective biases, after which a convolution of the Gaussian function with the dirty beam is subtracted from the dirty image, and the Gaussian is added to the model image. The conversion from the model image to the clean image remains unchanged, but should yield a better result as the model image is of better quality.

The clean image should have a lot less sidelobe structure than the dirty image, allowing for a more accurate study of the source. As said before, the model produced by the CLEAN process can be used for further self-calibration of the data, giving even more accurate results. This does require to run CLEAN after each self-calibration round, so it is important to check the amount of quality gained after each calibration-CLEAN iteration, as improvements in image quality for the used computing time will go down for each iteration.

3 Observations and Data Processing

3.1 LOFAR

The observations for this thesis were done using the Low-Frequency Array (LOFAR). This section gives a short explanation of its working. An overview is given of the distribution of its stations, followed by a description of the design of the antennas. After that, the layout of the stations and the placement of the antennas in the stations is explained. To round off, a short summary of the data processing path is given. For a more detailed description of LOFAR, we refer to van Haarlem et al. (2013).

3.1.1 LOFAR overview

The LOw-Frequency ARray (LOFAR) is a radio interferometer with stations spread over Europe, with its core located in the north of the Netherlands. Initially, 40 LOFAR stations were built in the Netherlands, 5 in Germany, and 1 in France, Sweden and the UK each. By the time of writing, 1 more station has been added in Germany, Ireland and Latvia each and 3 in Poland, summing up to a total of 54 stations spread over Europe, which together form the LOFAR International Telescope (ILT). These stations work together as a phased array: every antenna on its own can observe the whole sky, but the combined signal of the antennas can be pointed artificially by adding a delay to the signal of the respective antennas. Using this technique makes LOFAR a nimble telescope, as it can observe events and sources in almost any preferred direction. At the same time it also makes LOFAR very precise, as the angular resolution of a combined array is much better than that of an individual antenna, as discussed in section 2.

Operating at frequencies between 10 and 240 MHz (wavelengths 30-1.2 m), long baselines are of great importance for LOFAR as they provide the ability to resolve objects and structures of the smallest angular sizes, even at these low observing frequencies. This, together with the wide range of baselines (68 meters in the core up to 2000 kilometers between international stations), means that many source structural sizes can be resolved simultaneously. This allows for precise imaging of complex target sources, which is crucial for the scientific goals that LOFAR was built for, which include:

- Studying the epoch of reionization by probing the 21-cm line between redshift 11.4 and 6, to gain insight in the process of ionization of the intergalactic medium by the first stars and black holes.
- Searching the radio sky for high redshift radio galaxies and clusters to understand their morphology and their contribution to the cosmic star formation history.
- Observing a variety of "short-term" events in the radio sky. These include planetary and stellar flares on the timescale of microseconds to hours, but also small AGN outbursts, supernovae and gamma-ray bursts which can last several years. Solar flares in particular are more widely studied than just in astrophysics, because they strongly influence the Earth's environment.
- Studying the acceleration mechanisms responsible for the production of high-energy cosmic rays. These mechanisms are often the result of strong magnetic fields and shock waves, which are seen in radio galaxies as well as in supernovae.

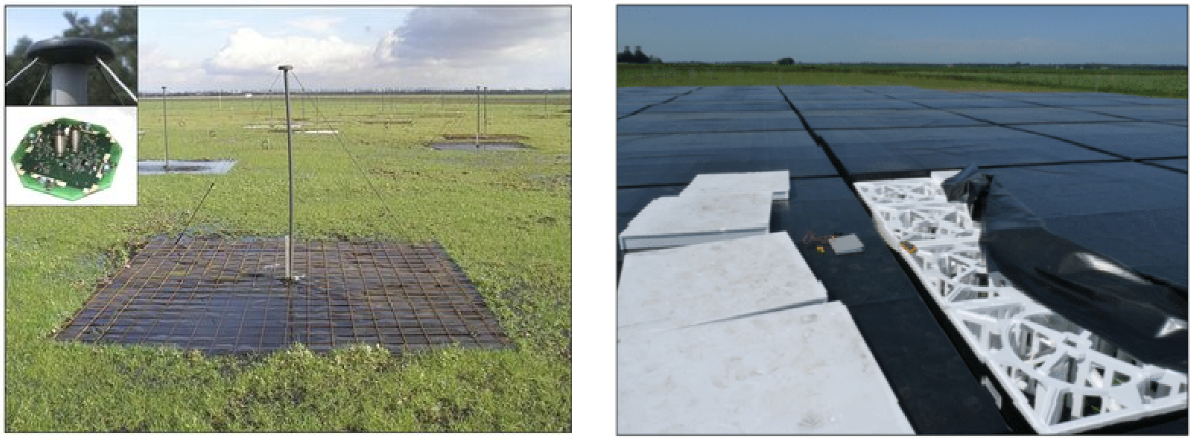

Figure 10: **Left:** Image of a LOFAR LBA dipole antenna, showing the ground plane, the vertical PVC tube with its molded cap, and the 4 copper wires. **Right:** Image of a LOFAR HBA, with part of the protective cover taken off to show the antenna architecture. (van Haarlem et al., 2013)

3.1.2 Antenna design

The wide range of baselines and the large number of required antennas brought a few consequences concerning the construction of the antennas so that mass production was feasible: it needed to be simple and low-priced. However, they also needed to be able to withstand the sometimes violent weather conditions in order to stay operative for more than just a few years. The resulting design is different from the traditional dish interferometer. Figure 10 shows images of the two chosen antenna types: the Low-Band Array (LBA) and the High-Band Array (HBA).

The LBA is the part of the array observing the frequency range between the low-frequency end of the radio window at 10 MHz, and the FM radio band at 90 MHz. In practice, this is reduced to a range between 30 and 80 MHz due to ionospheric reflection of radio waves below 30 MHz and the effect of the FM band upward from 80 MHz. An LBA antenna consists of 4 copper wires that together observe two orthogonal linear polarizations. These wires are attached to a ground plane via a rubber spring on one end, and to a molded cap on top of a vertical PVC tube on the other end. This molded cap contains a low-noise amplifier (LNA), and two coaxial cables running down through the PVC tube to transport both the signal output of the two polarizations and the power for the LNA.

The HBA is the part of the telescope that is being used in this thesis. It observes frequencies between 110 and 250 MHz, which becomes 110-240 MHz because of RFI above 240 MHz. Also contained in this band is the frequency used by P2000, a network used by the Dutch emergency services, which causes an RFI spike around 170 MHz. In the frequency range observed by the HBA, the electronics produce a significant portion of the observed noise, hence the design of the HBA is different than that of the LBA. HBA antennas are paired in a square of 4 by 4, which together form a ”tile”. The 16 signals from the antennas are then added in phase, dependent on the observational direction, forming a single ”tile beam”. Each HBA tile has its own analog electronics to delay and sum the individual antenna signals, meaning that every single tile essentially operates

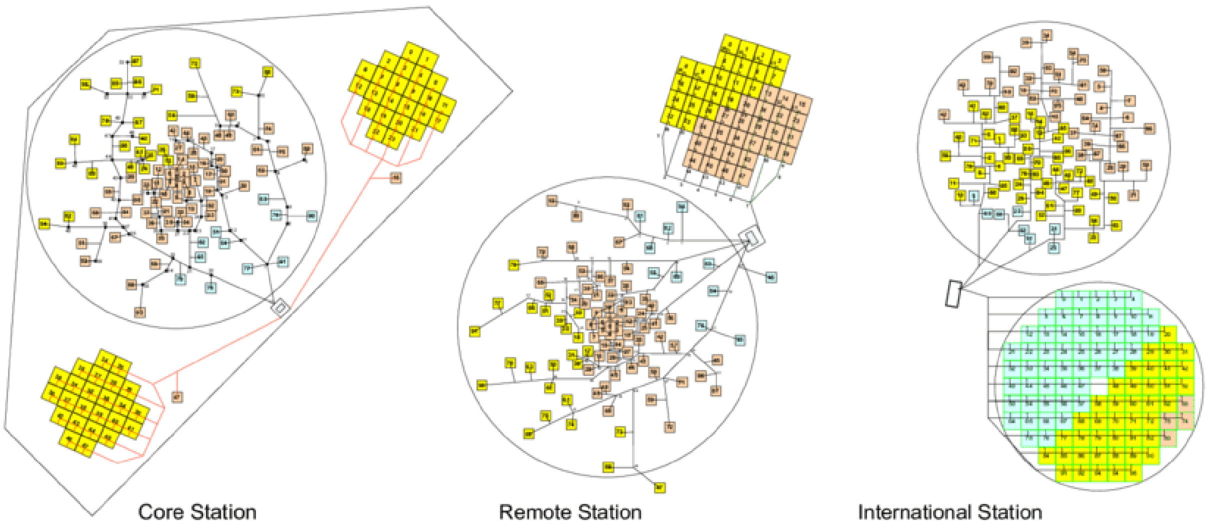

Figure 11: Layout of the core, remote and international stations. In each station, the LBA antennas are encircled, and the HBA tiles are represented by the larger closely packed squares. The difference in HBA tile count between the station types can be seen. It can also be seen that the LBA antennas in the core and remote stations are spread pseudorandomly, as opposed to the uniform spread in the international stations. All antennas are connected to a station cabinet that contains the electronics for the beamforming and the digitization and processing of the incoming signals. (van Haarlem et al., 2013)

as a phased array on its own. The antennas of the HBA are made from aluminium, supported by a ground plane and a polystyrene framework, and covered by two sheets of polypropylene foil.

3.1.3 Station layouts

Figure 11 shows the layout of the different LOFAR station types: core, remote and international. The core is located a few kilometers north of the village of Exloo. The nucleus of the core is a circle-shaped island hosting 6 of the core stations, the "Superterp", seen in Figure 12. An additional 18 core stations can be found within 2 kilometers from the Superterp. Further away from the core we find the remote stations. These are scattered through the north of the Netherlands, with baselines reaching up to 120 kilometers. A further 13 international stations lie scattered over Europe. The international stations were not used in this research, due to calibration issues. This caused their contribution to worsen the image quality to a level where the loss of angular resolution by excluding them was less harmful.

A few key differences can be seen in Figure 11. The biggest difference is seen in the location of the HBA tiles, and the number per station. It is seen that the core stations have two groups of 24 tiles positioned on opposing sides of the LBA stations in a "Mickey Mouse ears" configuration, while remote stations have one group of 48 tiles and international stations have 96 tiles. The spacing between the core station tiles has the big advantage that it can be used as a short baseline, even shorter than the baselines between two stations and thus sensitive to larger angular sizes. These

Figure 12: The LOFAR Superterp, the central part of the LOFAR core. The Superterp contains 6 core stations, each consisting of 96 low-band antennas and two groups of 24 high-band antenna tiles. In the bottom left and top right, 3 more core stations can be seen. (van Haarlem et al., 2013)

small spacings can also be used to form a more uniform *uv*-coverage at small *uv*-distances. For the remote stations and the international stations this is not as much of an advantage, as such a configuration would form near-duplicate baselines. Instead, their tiles are grouped all together to benefit from better sensitivity, which is a big advantage on the long baselines.

It can also be seen that the LBA antennas of the core and remote stations are distributed pseudorandomly to form as many different baselines as possible, which is advantageous for the uv-coverage. The antennas are distributed over a circular area, with the inner part of the circle being more densely packed than the outer circle. A distinction is made between the ”LBA Inner” and the ”LBA Outer” configurations, each consisting of 48 antennas. Both parts have their advantages when it comes to field of view and sensitivity, which is why an observer can choose one configuration that suits their observation best. The LBA antennas in the international stations are distributed uniformly.

3.1.4 Further processing

In Figure 11, it can also be seen that all antennas are connected to a station cabinet. This cabinet contains the computer responsible for the beamforming, i.e. the calculation of the signal delays for every antenna given a certain pointing direction, as well as the receiver unit and the electronics responsible for the further processing of the received signals. The analogue signals from the antennas are filtered, amplified and digitized in the receiver unit. After the signal delays have been applied in the beamforming step, the signal is streamed to the central processing facility located at the Center of Information Technology at the Zernike Campus in Groningen. An IBM Blue Gene/P supercomputer was originally used to either process the data-stream in real-time, or to store the data packs for correlation later. This task has now been transferred to COBALT, a GPU based supercomputer. After being processed by the cluster, the data packs are written to the storage cluster, where a range of pipelines can be used to reduce the data packs so that they can be stored on the LOFAR long-term archive (LTA) for scientific purposes.

3.2 Observations

The observational data of Cygnus A obtained by LOFAR has been processed to obtain images needed for this research. This section describes the steps that have been taken in calibrating and imaging the observations to obtain images and spectra suitable for scientific analysis, and gives a justification of certain choices that have been made during this process.

3.2.1 Data retrieval, description & processing

The HBA observations of Cygnus A were obtained from the LTA, and stored on a hard disk of the Kapteyn Astronomical Institute. Two data packs were obtained, one in the 210-250 MHz range and the other in the 110-190 MHz range. The data sets have a size of 1 TB and 2 TB respectively, adding up to 3 TB in total. These data sets in turn consist of MeasurementSet directories (MS) of 5 GB, each covering a frequency band of width 192 kHz.

The data in the frequency range 110-190 MHz, from now on referred to as the ”mid band”, was observed on May 13, 2014 between 00:11 and 12:11 UTC. The array configuration consisted of 23 core stations, 13 remote stations and 8 international stations. The spread of the core and remote stations can be seen in Figure 13 and Figure 14. The data were taken using the 200 MHz sampling clock i.e. the observation lies in the second Nyquist zone. Measured visibilities were averaged every 2.01 s. The 399 produced MS each consist of 4 polarizations and 4 spectral channels of width 48.828 kHz. Hence, the observation covers 77.92 MHz of the 80 MHz bandwidth.

The 210-250 MHz data, from now on referred to as the ”high band”, was observed on May 14, 2014 from 00:00 to 12:00 UTC. 203 MS were produced, also consisting of 4 polarizations and 4 spectral channels with a bandwidth of 48.828 kHz each. The observation therefore covers 39.65 MHz of the total 40 MHz bandwidth. Visibilities were again averaged every 2.01 s using the 200 MHz sampling clock, meaning that the data were located in the third Nyquist sampling zone.

The further processing needed to convert the data into scientifically usable images was done using the *Common Astronomy Software Applications* (CASA), version 4.5.0. At the time of the processing this version had been superseded, but the age of the MS created conflicts with the newest versions of CASA. The used version contained all functionalities necessary to obtain the results desired in this research.

3.2.2 Splitting subbands

Each subband contained 5 GB of data, so it was advantageous to average the data, to speed up the calibration and imaging process. For each subband, the 4 spectral channels were averaged into 1 channel. The polarizations XY and YX were removed, keeping only the XX and YY polarizations. The data from the international stations were also removed, because using them was not beneficial for the image quality. The remaining subband size was 1.5 GB, which significantly sped up the calibration.

3.2.3 Calibration & Imaging

The resulting data packages still needed to be calibrated. This was done by using the task GAINCAL in CASA, which calibrates the antenna gains as a function of time to remove atmospheric and instrumental errors, as described in Section 2.3. First, 37 subbands of the mid band data

Figure 13: Distribution of the core stations used during the observation. The map covers an area of approximately 10 km^2. The Superterp, located 3.5 kilometers north of the village of Exloo, is the circle shaped island located in the center of the image. Station CS006, located on the Superterp, was not active in the observation.

Figure 14: Distribution of the remote stations used for the observation. The map spans approximately 13 000 km^2 in area. For reference, the islands Ameland and Schiermonnikoog can be seen at the top of the image, and the Noordoostpolder to the west of station RS310. Also labeled is the Astron facility near the town of Dwingeloo.

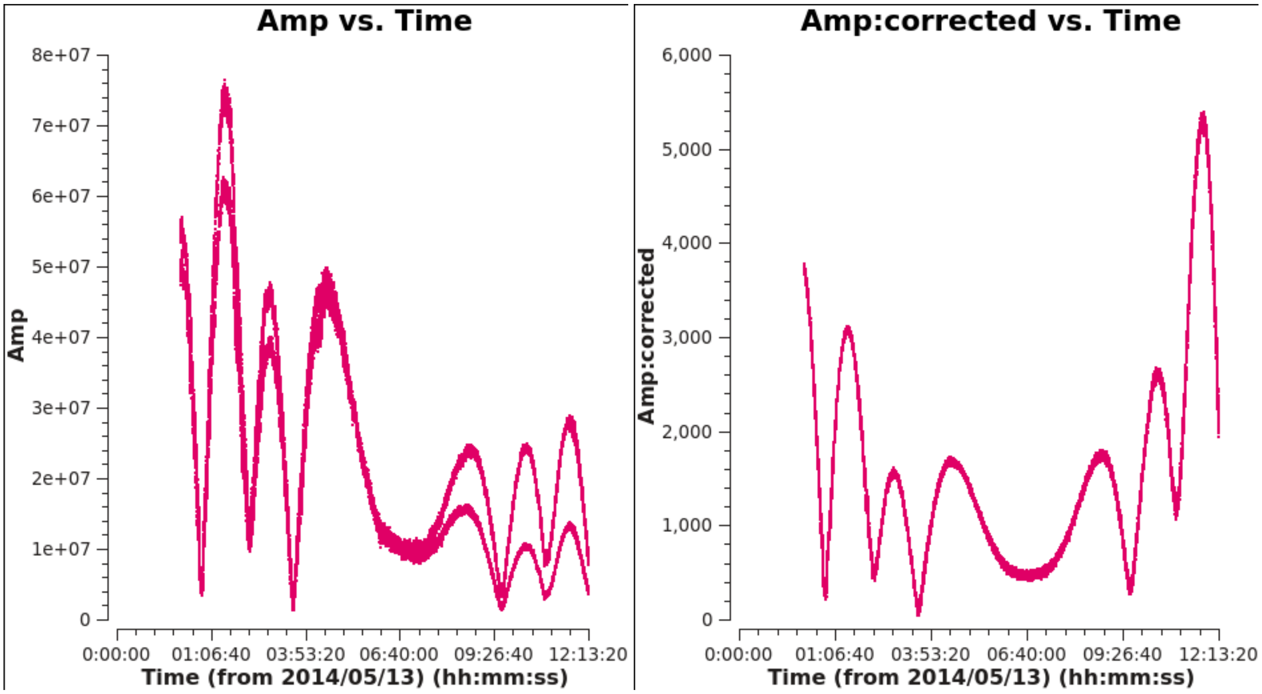

Figure 15: Before-and-after of the calibration process on the visibility amplitudes of baseline CS002&RS406, in the HBA mid band, centered at 150 MHz. These gain tables were calculated by the GAINCAL tool in CASA. The horizontal axes are the time of the observation in UTC. **Left:** Uncalibrated visibility amplitudes. Their values are seen to be orders of magnitude higher than Cygnus A's flux density at 150 MHz, and many fluctuations are seen on small timescales. In addition, the amplitudes differ between the two polarizations. **Right:** Visibility amplitudes after application of the calculated amplitude gain table. The corrected amplitude values are now in the expected regime, and changes are more continuous over time. The two polarizations are have the same amplitude as well.

set spread across the total frequency range were calibrated to obtain a rough model of the data between 110 and 180 MHz. This calibration was done using a model of Cygnus A provided by McKean et al. (2016), which consisted of two Taylor terms. The first Taylor term can be interpreted as the intensity at the central frequency of the band, and the second term tells how the flux changes as a function of frequency, assuming a power law behaviour as derived in Section 1.3.1. An image showing the spectral index can then produced from this second Taylor term image by CASA. An example of the result of the self-calibration using this model on the amplitudes and phases in one band can be seen in Figures 15 and 16. The corresponding gain tables can be seen in Figure 17.

After self-calibration of the individual bands, 37 images and models were produced by the multi-scale CLEAN algorithm (using scales of 0, 5, 10 and 20 arcseconds) with superuniform weighting. The fluxes of the resulting models, each consisting of one Taylor term, were found to be slightly off (up to 2%) from their theoretically expected fluxes, given by a third order polynomial

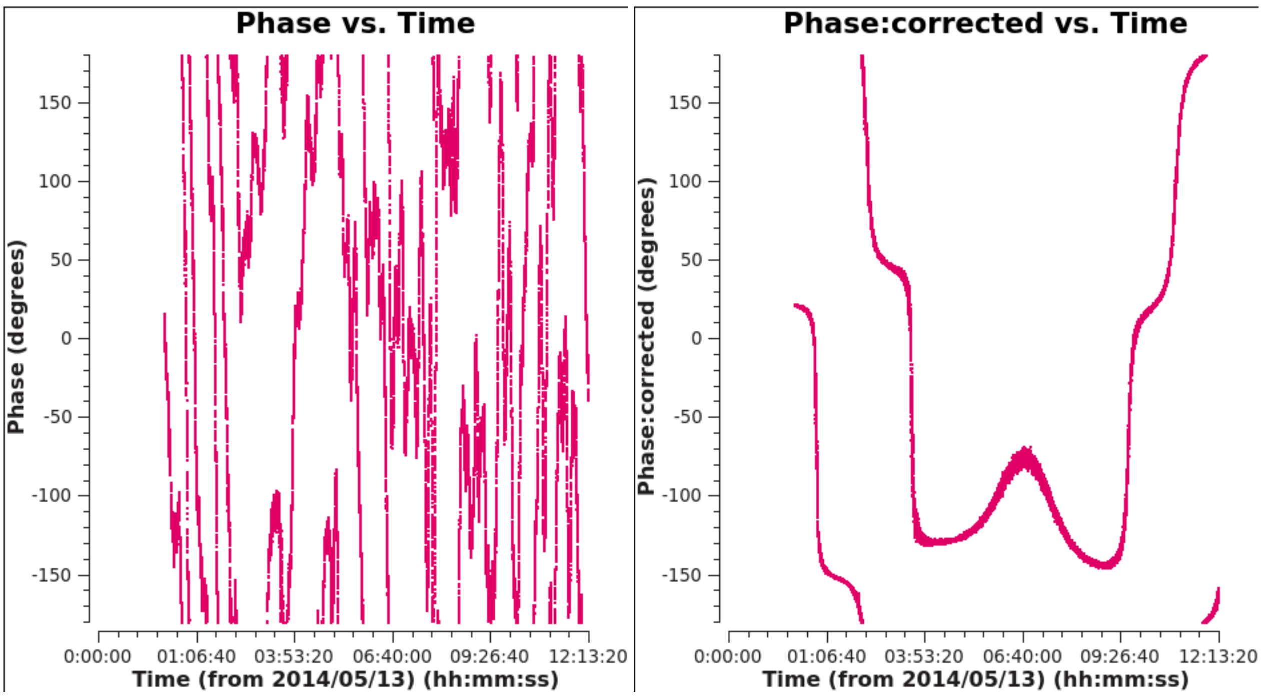

Figure 16: Before-and-after of the calibration process on the visibility phases of baseline CS002&RS406, in the HBA mid band, centered at 150 MHz. These gain tables were calculated by the GAINCAL tool in CASA. The horizontal axes are the time of the observation in UTC. **Left:** Visibility phases before calibration. Like the amplitudes, the phases show many small timescale fluctuations, as well as rapid changing over time, up to two revolutions per hour. **Right:** Visibility phases after applying the calculated phase corrections. Similar to the amplitudes, less small timescale fluctuations are seen, and more reasonable phase changes per hour are seen.

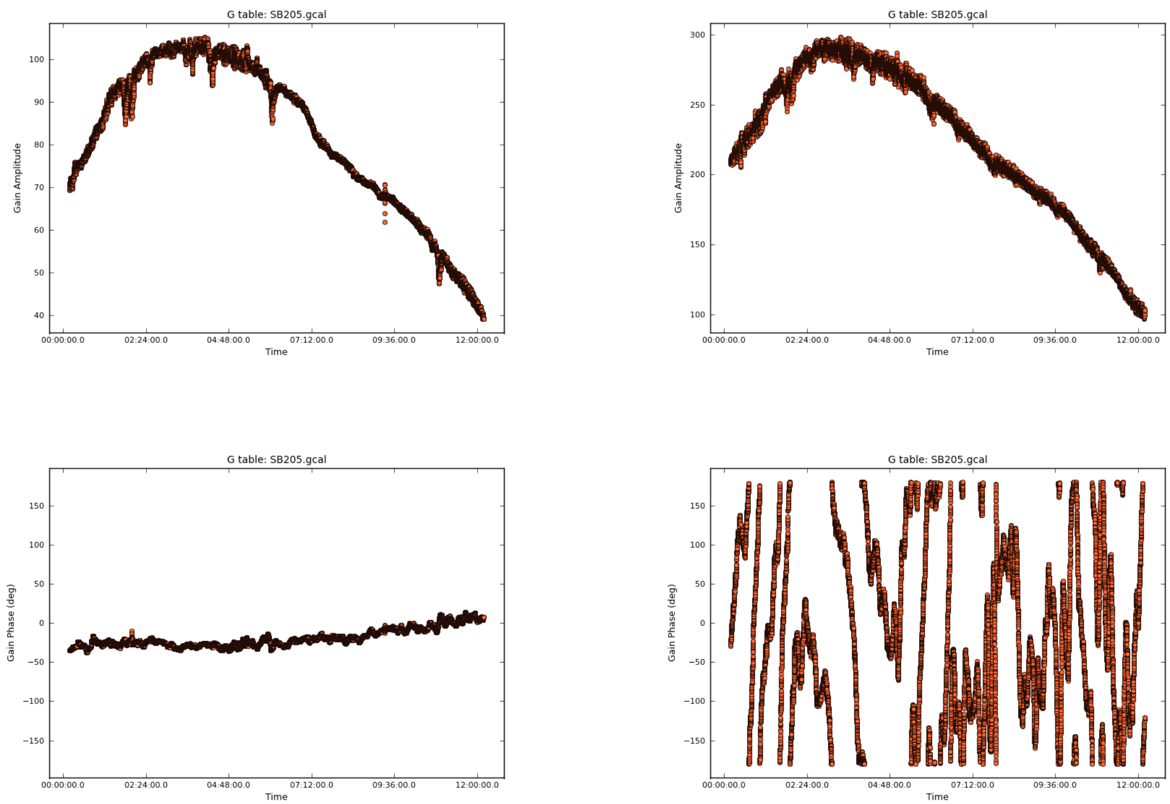

Figure 17: Gain calibration tables of the observation shown in Figures 15 and 16. **Top left:** Amplitude calibration of CS002. The overall appearance is the consequence of Cygnus A's movement over the sky, and small fluctuations are the corrections of atmospheric and electronic changes. **Top right:** Amplitude calibration of RS406. The overall appearance is the same, but small timescale variations (instrumental corrections) are different from CS002. **Bottom left:** Phase calibration of CS002. These corrections are relatively small, as CS002 lies close to the reference antenna CS001, and thus the geometric delay does not change a lot over time. **Bottom right:** Phase calibration of RS406. This antenna lies relatively far from the reference antenna, causing the geometric delay to change more rapidly over time, and hence larger corrections are necessary.

of the form

$$S[\mathrm{Jy}] = A_0 \prod_{i=1}^{3} 10^{A_i \log_{10}^{i}(\nu/150\,\mathrm{MHz})}, \tag{30}$$

where the values of the constants $A_0 = 10690 \pm 230$, $A_1 = -0.67 \pm 0.19$, $A_2 = -0.240 \pm 0.017$ and $A_3 = 0.021 \pm 0.014$ were found by McKean et al. (2016). After scaling the models to the correct fluxes and another round of self-calibration, these bands were then concatenated and imaged using the Multi-Scale Multi-Frequency-Synthesis (MS-MFS) CLEAN algoritm. The resulting model with two Taylor terms and superuniform weighting was then used to calibrate the complete mid band data set, leaving out several subbands (mostly near the edge of the band) because of RFI. Using the same approach, all subbands underwent a flux correction and an extra round of self-calibration. The resulting calibrated subbands were concatenated into one MeasurementSet, which underwent a few iterations of self-calibration. To remove some artifacts caused by the hotspots, multiple weightings (Briggs, uniform and superuniform) were used, resulting in a noise reduction and thereby a higher dynamic range.

The resulting 2 Taylor term model was then used to calibrate the data set between 210 and 250 MHz (the "high band", consisting of 203 subbands). This was done by first calibrating the middle 40 subbands in the same fashion as the mid band, and then using the models of the outer bands to calibrate outward in steps of 40 subbands. Similar to the mid band data set, all subbands were concatenated (excluding the ones containing RFI), followed by a few extra rounds of self-calibration and imaging.

Merging the mid and high band, a measurement set containing all the useful data of the LOFAR HBA was obtained. The *uv*-coverage can be seen in Figure 18. The full measurement set underwent a few extra rounds of self-calibration and flagging to further improve the image quality. The final imaging was done using 3 Taylor terms, where the third Taylor term represents the quadratic change in flux density as a function of frequency. The fitted spectral behaviour in this case follows a power law of the form

$$S_\nu = S_{\nu_0} \left(\frac{\nu}{\nu_0}\right)^{\alpha + \beta \log(\nu/\nu_0)}, \tag{31}$$

where β is the spectral curvature. This equation is very similar to the power law derived in Equation (15), but the addition of $\beta \log(\nu/\nu_0)$ allows the spectrum to show curvature in logarithmic space. Analogous to the spectral index image, a spectral curvature map is calculated from the third Taylor term image by CASA.

3.2.4 Obtaining spectra

For the spectral analysis of the images, it was necessary to image all frequency bands separately, in order not to lose any spectral behaviour that is being averaged out by the MFS imaging, such as spectral turnovers and aging breaks. This allows for the fitting of models of the absorption processes explained in Section 1.3.2 to the data to get a more complete understanding of the physical properties of certain regions and the resulting contribution of the absorption processes. The imaging was done in the same way as the MFS images, using the same deconvolution scales but with only one Taylor term. A restored intensity image was obtained for all 492 frequency bands.

In hotspots A and D, the LOFAR HBA observing frequencies only cover the frequency range near the spectral turnover point. Adding data points outside this range allows for a more precise

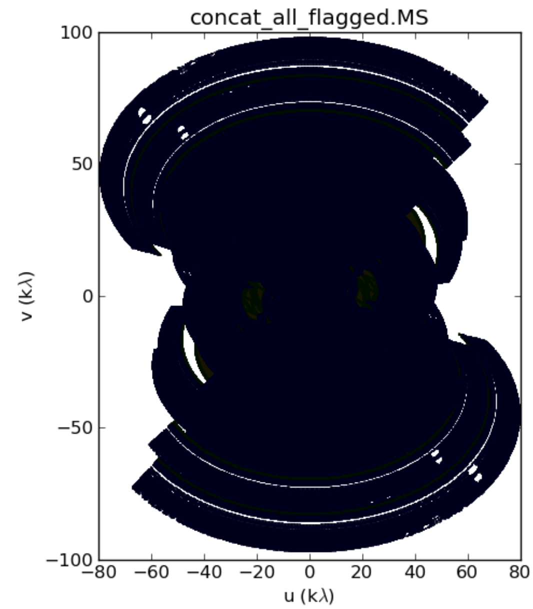

Figure 18: *uv*-Coverage of the final data set, plotted on wavelength axes. Due to the wavelength range of the data set, the arcs traced out by the baselines are smeared out radially, allowing the *uv*-plane to be filled quite well. Some small gaps still arise from removed bad data.

fitting of the absorption models, as the optically thick and optically thin spectral slope are two important parameters in the fitting process. The used additional data points lie in the optically thin regime, and were observed by the VLA. Two L-band images at 1.037 and 1.424 GHz were obtained from the work of Arwa Dabbech (2021), and one P-band image at 327 MHz from the VLA + Pie Town link interferometer from Lazio et al. (2006). The addition of these images can provide a more robust fit of the models on the optically thin end compared to using the LOFAR data only.

Measuring the spectra comes at a cost, as all the LOFAR and VLA images have different restoring beams. To make a fair comparison between the spectral bands, they were all smooted to the same beam of 4.7×3.2 arcseconds, which is the beam size of the 110 MHz LOFAR band. This smoothing guarantees that beam smearing is the same over all bands, and thus all pixels experience the same overflow of surface brightness from other nearby pixels. The consequence is that the measured spectra of most features are flatter than their 'real' spectra, as a significant contribution from nearby radio lobe plasma is beam smeared into the features from which the spectra are taken. An alternative approach to measure the flux densities would be from the CLEAN components in the model images. This approach would also come at a cost, as the model images super-resolve the brightness distribution by collapsing structures of multiple pixels into single pixels, and therefore this strategy might not be ideal either. After the smoothing of all bands, spectra were obtained in a total of six features by measuring the maximum surface brightness pixel value inside a radius of 2.5 arcseconds around the region of interest. The uncertainty in each measurement is calculated as a percentage in the flux calibration error from Equation (30), added to the rms noise in the residual image of each band. The uncertainties are dominated by the flux calibration error, which vary between 2 and 5 percent.

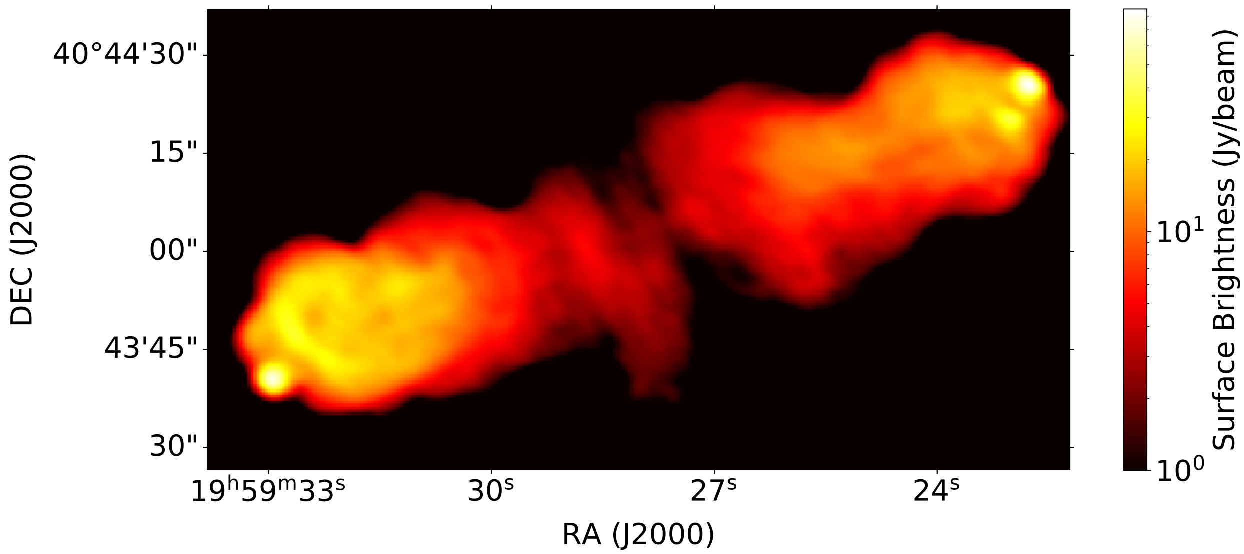

Figure 19: Resulting intensity map of Cygnus A between 110 and 250 MHz obtained using the MS-MFS CLEAN algorithm. This is the first of three Taylor terms used to describe the spectral behaviour. Four spatial scales were used for the deconvolution (0, 5, 10 and 20 arcseconds). The restoring beam size of the image has dimensions of 2.03 $\times$ 1.38 arcseconds, with an RMS noise of 16 mJy/beam and maximum intensity of 85 Jy/beam. The image has a dynamic range of 5300. A logarithmic scaling is chosen for the colormap not to lose detail in the less luminous regions of the lobes.

4 Results

This section will present the results of the data reduction. A description will be given of the images resulting from the self-calibration and CLEAN process, including an investigation of the features that can be seen in them. The obtained spectra will also be shown, as well as the emission models that have been fitted to them.

4.1 Imaging

The self-calibration of the full LOFAR MeasurementSet produced a dataset that could be imaged with an rms noise of 16 mJy/beam and an off-source noise to 30 mJy/beam. The restoring beam of the image is of size 2.03 $\times$ 1.38 arcseconds, which just about resolves hotspots A (lobe) and D (counter-lobe). Both the rms noise and the beam size are an improvement compared to the imaging by McKean et al. (2016), who obtained an rms noise of 43 mJy/beam and a beam size of 3.8 $\times$ 2.7 arcseconds. The highest surface brightness in the image is 85.1 Jy/beam, which compared to the rms noise gives a dynamic range of $\sim$ 5300. The resulting image, spectral index map, spectral index uncertainty map and spectral curvature map can be seen in Figures 19 to 22.

For comparison with Cygnus A's radio structure, an image taken with the Chandra X-ray Observatory was obtained from Sebokolodi et al. (2020), shown in Figure 23. Being in the 0.5-0.7 keV regime, the image shows a lot of additional structure of Cygnus A that is not seen in the radio observations. All hotspots can be identified in the same locations as in the radio map, but the rest

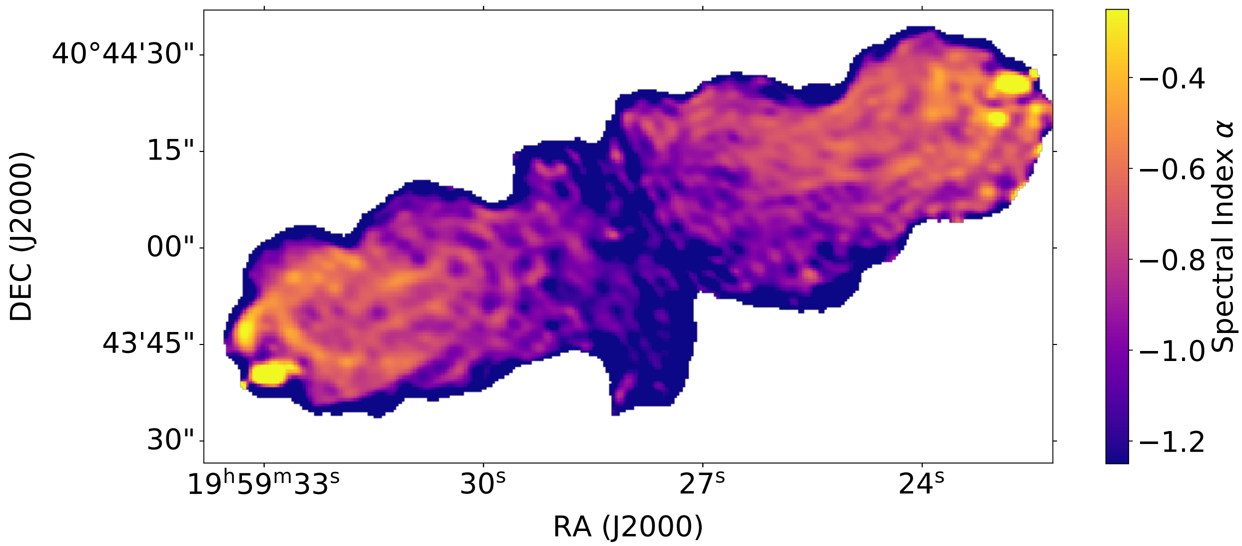

Figure 20: Spectral index map of Cygnus A, derived from the first and second Taylor term produced by MS-MFS CLEAN. Values of the spectral index in the hotspots reach as high as +0.3, but the scale is set between -0.25 and -1.25 to show more structure in the lobes. The spectral index in the lobes varies between -0.5 and -1, and appears to show a correlation between structures that have flat spectra and a high flux density.

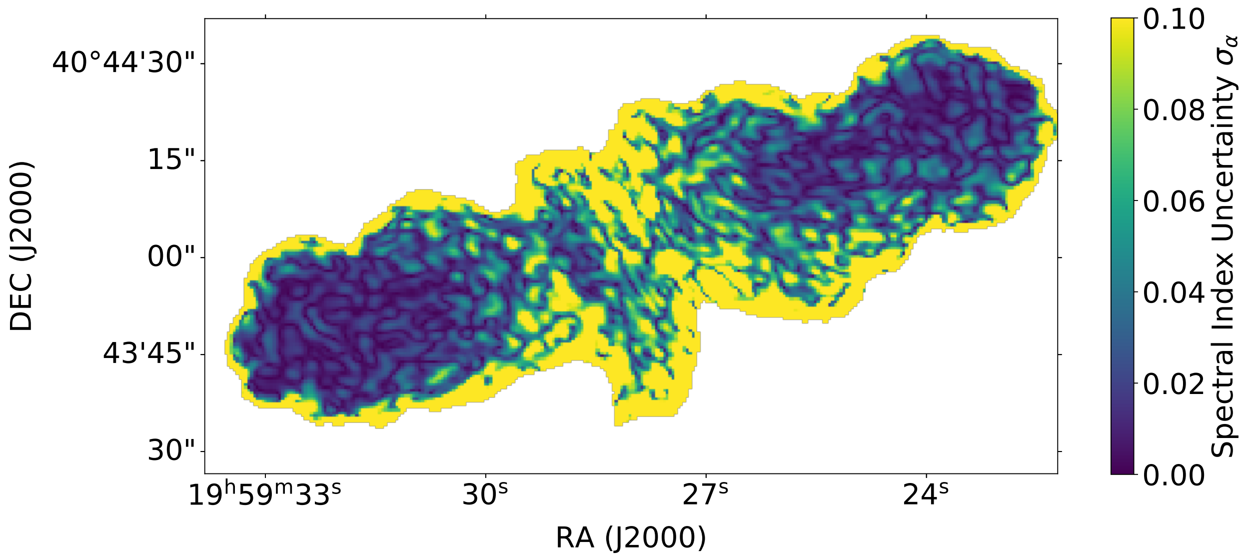

Figure 21: Uncertainty in the spectral index map of Cygnus A, calculated by MS-MFS CLEAN. The colormap is saturated at an error of 0.1, showing that the hotspots and the brightest parts of the lobes often have uncertainties smaller than 0.02. Uncertainties higher than 0.1 are seen near the edge of the mask and in other regions of low luminosity.

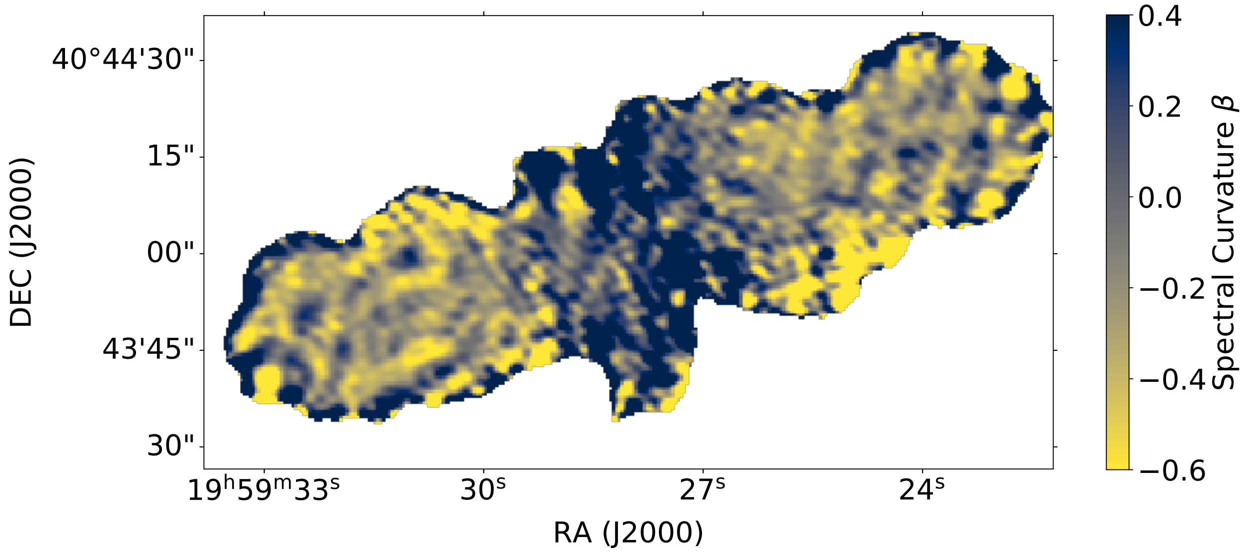

Figure 22: Spectral curvature map of Cygnus A, resulting from combining the first, second and third Taylor term image fitted by the MS-MFS CLEAN algorithm. The spectral curvature takes on extreme values in regions of low intensity, but setting proper boundaries on the color axis reveals the underlying structure that allows for the identification of the same lobe components seen in the intensity map and the spectral index map, such as the hotspots and other bright features.

of the lobe's internal structure is not seen. What can be seen are the X-ray jets, although they do not always coincide with the radio jets. The AGN itself can be seen as well, being the brightest region in the entire image. Perfectly fitting around the cavities filled by the radio lobes, an extended X-ray structure can be seen, known as the cocoon. This emission is likely coming from hot gas that is being heated, shocked and pushed outward by the expanding radio lobes.

A lot of complex structure in Cygnus A's lobes can be seen in the radio images. The overall structure in the intensity map shows the key components that have been observed in earlier radio observations from 150 MHz to 22 GHz by Leahy et al. (1989) and Carilli et al. (1991). In the lobe, which is located westwards of the center, the intensity map clearly shows both secondary hotspot A (most northern) and primary hotspot B, accompanied by brightly emitting lobe plasma. The rest of the lobe shows to have more dimly emitting structure towards the center, where even a jet feature can be seen swerving towards hotspot B. The spectral index map shows that the same jet component has a slightly flatter spectrum than the surrounding lobe plasma. Really standing out in the northwest of the spectral index map are the hotspots, where hotspot A peaks at $\alpha = 0.22 \pm 0.02$ and hotspot B at $\alpha = -0.02 \pm 0.02$. They also show a strong negative curvature, making them easily identifiable in the curvature map, contrary to the jet, which is not identified in the curvature map.

The central black hole can not be distinguished in the intensity map, while it can be seen sitting in the center of the spectral index map, having a significantly flatter spectrum than its surroundings. What can also be seen at the center is the plume, the third lobe of radio emitting plasma located southward of the black hole. This emission is expected to be old, as it is thought to be the result of earlier jet activity, meaning that this region in the spectral index map should show a very steep slope and curvature due to aging. While this is the case, this has to be taken with a grain of salt

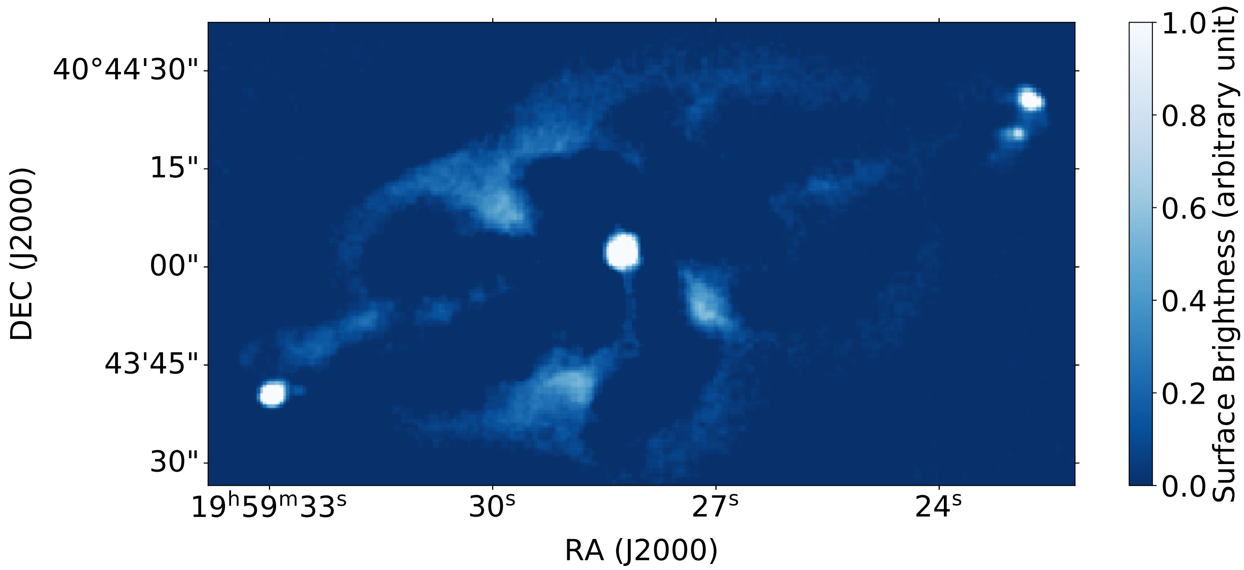

Figure 23: Chandra X-ray image of Cygnus A at 0.5-0.7 keV, adapted from Sebokolodi et al. (2020) covering the same sky coordinates as the radio images. Scaling and normalization is chosen such that bright extended structures are seen clearly. The nucleus is seen in the center, and the hotspots in the top right and bottom left corner. Also seen is the expanding X-ray jet and the cocoon draped around the radio lobes.

as the observed emission in this region is relatively faint.

The counter-lobe, which is found to the east of the center, is brighter than the lobe. Secondary hotspot D can clearly be seen in the most southeast corner of the counter-lobe, while primary hotspot E is more difficult to find in the intensity map because it is blent with hotspot D. A look at the spectral index map shows its location more clearly, just to the northwest of hotspot D. The spectral index peaks at $\alpha = 0.48 \pm 0.02$ for hotspot D and $\alpha = -0.05 \pm 0.02$ for hotspot E in their respective centers. The curvature map also clearly shows the location of hotspot D, while hotspot E is less dominantly present than in the spectral index map.

The three maps show a range of additional emission features in the counter-lobe. The first and least bright one is a pair of filamentary features lying parallel on both sides of the radio counter-jet in the center of the counter-lobe. Both features show to be brighter than their surroundings, the northern one being the brightest of the two. The spectral index and curvature map also show both features, where the spectrum of the northern one is flatter, but also curved more strongly. Both filaments have been observed by Perley et al. (1984), who propose that they might be the result of shocking by the counter-jet that moves between them at supersonic speeds. The radio and X-ray maps seem to be consistent with this theory, as the gap between the filaments nearly perfectly coincides with the X-ray counter-jet.

The largest and brightest feature is the extended semicircle just inward relative to hotspots D and E, nearly stretching across the full width of the counter-lobe. The intensity map shows that this is the brightest feature in Cygnus A at these wavelengths, excluding the four main hotspots. The spectral index and curvature map also show that the spectrum of the semicircle is relatively flat and strongly curved, especially towards the north. It can also be seen in the X-ray image that the X-ray counter-jet coincides with this feature, as it widens at the point where it enters the semicircle region. The spectral maps also show that the northern part of the structure extends into a faint blob with a strong curvature and a flat slope, behaviour similar to that of a hotspot. This blob shows up faintly in the X-ray map as well, opening the possibility that this is the relic of an abandoned hotspot.

A series of contour images can be seen in Figures 24 to 26. The contour overlays give a more direct comparison between the features seen in Figures 19, 20 and 23. The radio intensity overlay on the spectral index map (Figure 24) shows that generally high luminosity and flat spectra go hand in hand throughout the source, where the hotspots are the extreme cases being significantly brighter and flatter than the rest of the lobes. This implies that there are many sites of particle acceleration throughout the lobe. In Figures 25 and 26 the coincidence between the radio and X-ray features can be seen more directly. The jet can be seen penetrating the lobe between the filaments, and widening when it reaches the semicircle. The cocoon is seen surrounding the inner part of the lobe as well.

4.2 Spectral analysis

Seven spectra were obtained from the images, of which the locations can be seen in Figure 27. Each of the regions shows noteworthy behaviour in the radio images and spectral maps. The main focus lies on hotspot A and D, as they are the brightest features in the radio map. These have been shown by McKean et al. (2016) to contain a turnover in their spectrum, and are thus candidates for FFA and SSA modelling. The spectra of hotspot A and D can be seen in Figure 28 and Figure 29, respectively.

Primary hotspot B is clearly distinguishable from its secondary companion hotspot A in the radio intensity map, contrary to hotspot D and E who are not fully separated. The spectral maps and the Chandra image also show hotspot B as a source that stands out from the rest of the lobe. For this reason, a spectrum from hotspot B is also taken, as the curvature suggests that it is likely to show flattening near 110 MHz and possibly a turnover. Being a confirmed hotspot, this region is also suitable for SSA and FFA fitting. The spectrum of hotspot B can be seen in Figure 30.

A feature that has not been identified as a hotspot before is hotspot C, the blob located to the northeast of hotspot D. The spectral maps show that this region has both a positively sloped spectrum and a strong negative curvature, the behaviour also seen in the main hotspots. The possibility exists that this is a relic hotspot, resulting from earlier jet activity. This could fit within the ”denstist drill model” proposed by Scheuer (1982) and the precession of the jet observed by Steenbrugge and Blundell (2008), as this would constantly move the ”current” hotspot to a new location, leaving behind aging hotspot plasma. This would also explain the faintness and the size of hotspot C, as relic hotspots are expected to expand and cool over time. Its spectrum will therefore also be investigated, which is shown in Figure 31.

From the shape and location of the semicircle shaped region (from now on named R1) and the filaments parallel to the X-ray jet (from now on R2 (north) and R3 (south)) it can be seen that they have a different physical origin than the hotspots, but all of their spectra can be seen to be significantly brighter, flatter and more strongly curved than the rest of the counter-lobe. These characteristics indicate that the electrons in these regions gain in energy via some process, but it is unlikely that this energy comes from direct injection by the jet, as is the case of the hotspots. The spectra of these features will be investigated to see whether the curvature leads

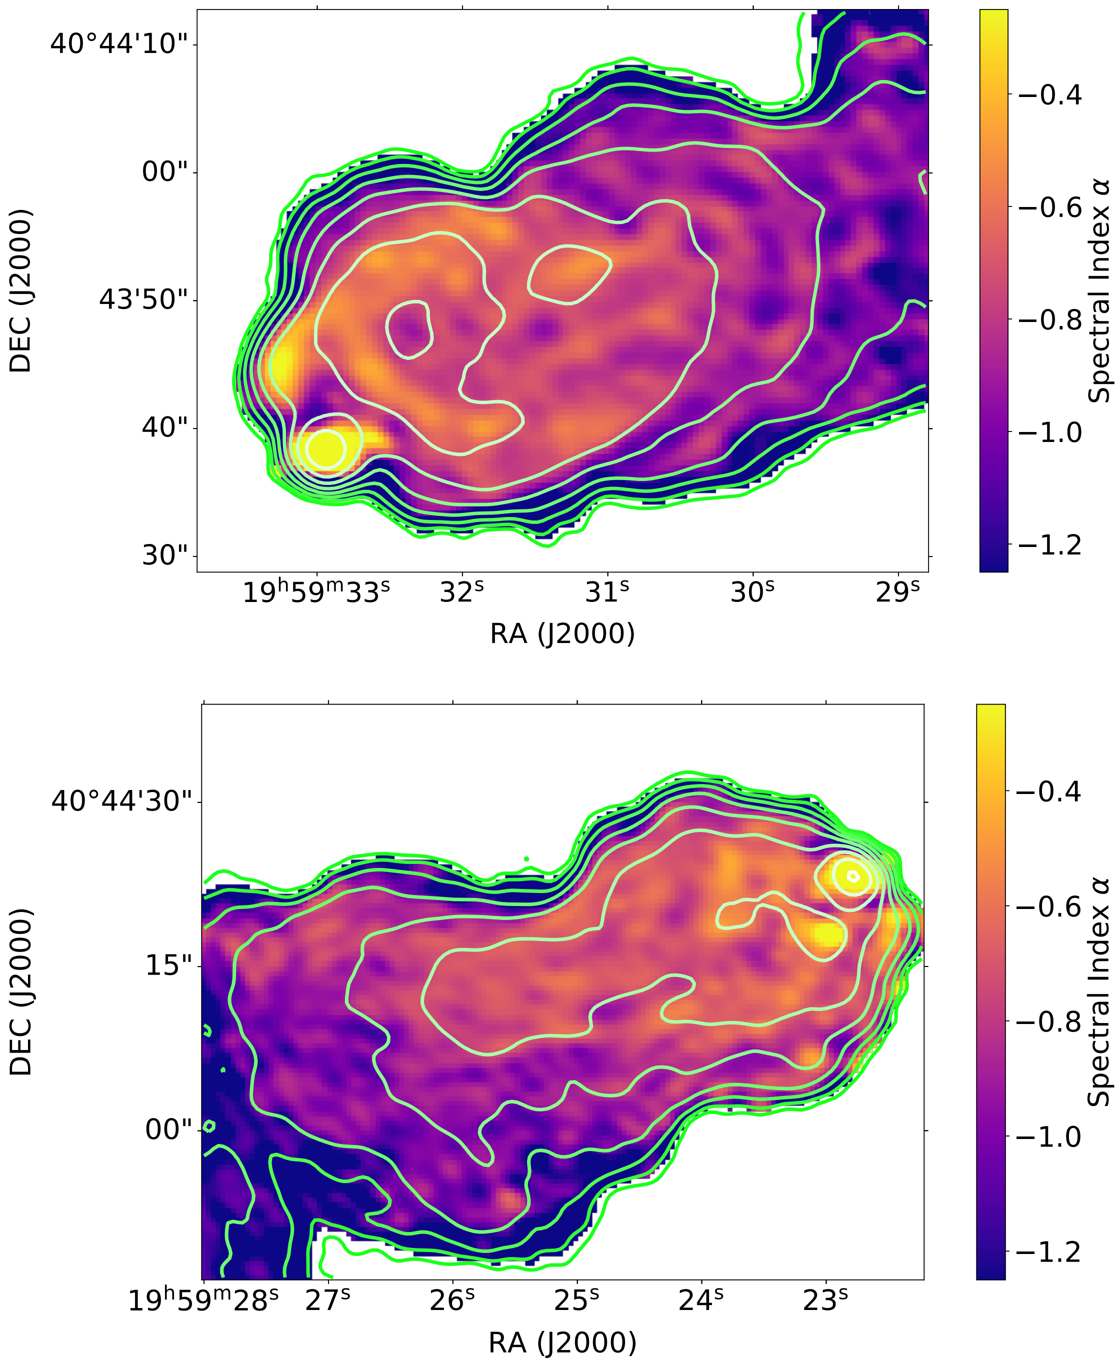

Figure 24: 110-250 MHz spectral index map with the total intensity contours overlayed of the counter-lobe (upper) and the lobe (lower). The range of the spectral index is set equal to the range in Figure 20 to show the same amount of detail in the lobes. The LOFAR contour values increase by factor of 2 in surface brightness, starting at 0.3 Jy/beam to clearly show the structure in both lobes.

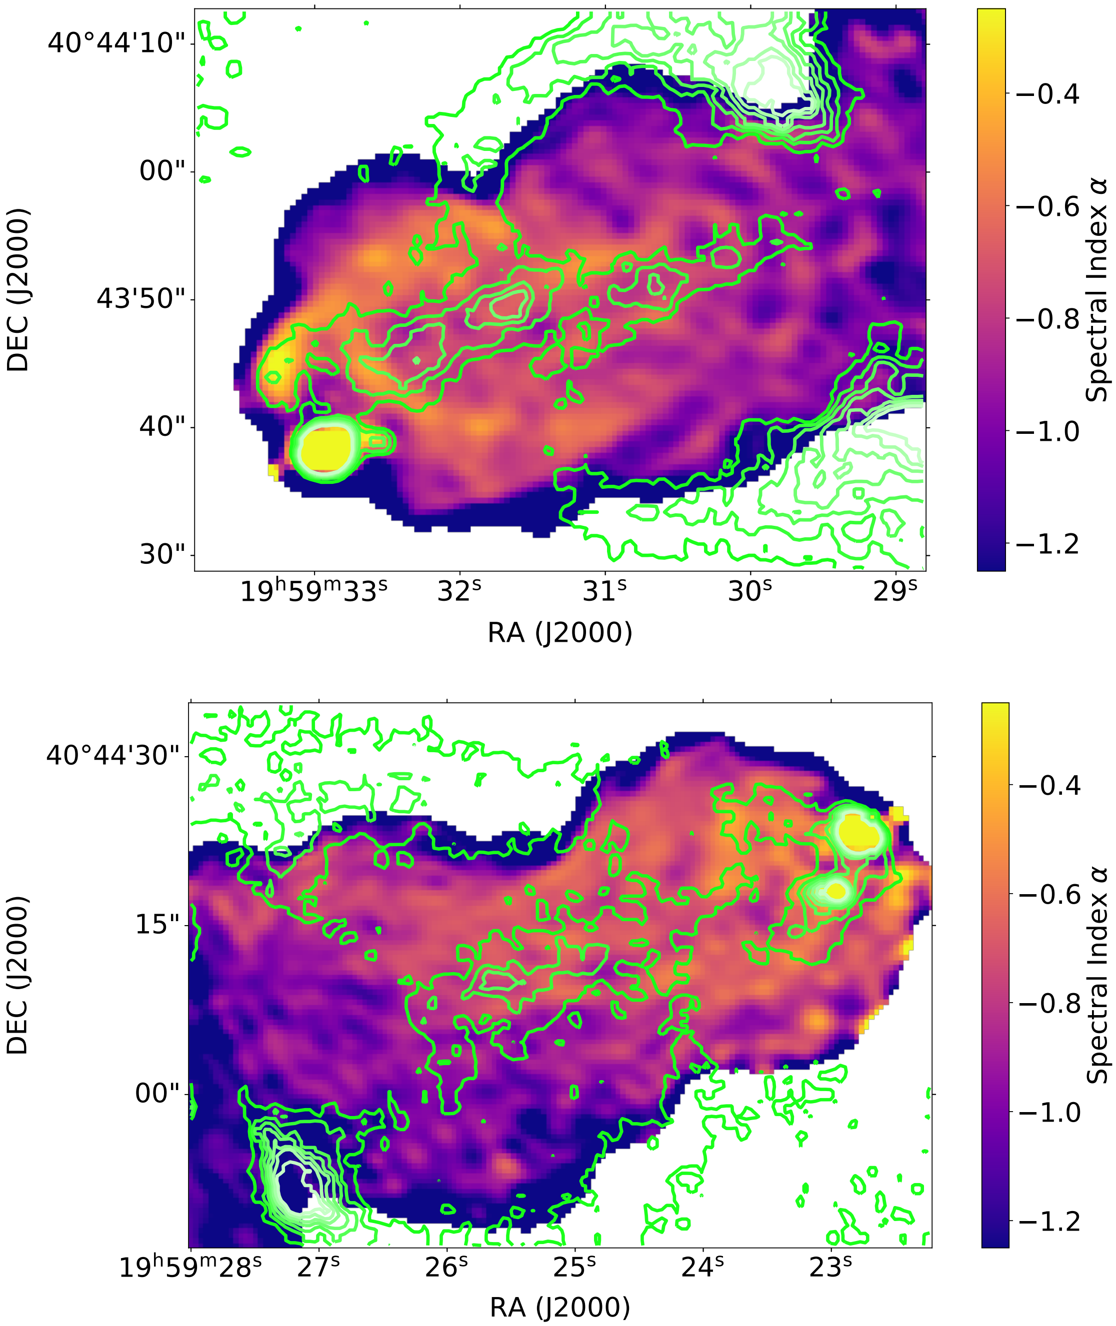

Figure 25: 110-250 MHz spectral index map with the 0.5-0.7 keV Chandra X-ray contours overlayed of the counter-lobe (upper) and the lobe (lower). A linear scaling in the intensity range between 1×10^{-9} and 2×10^{-7} is chosen for the X-ray contour levels to show the structure of the X-ray jet, the hotspots and the cocoon.

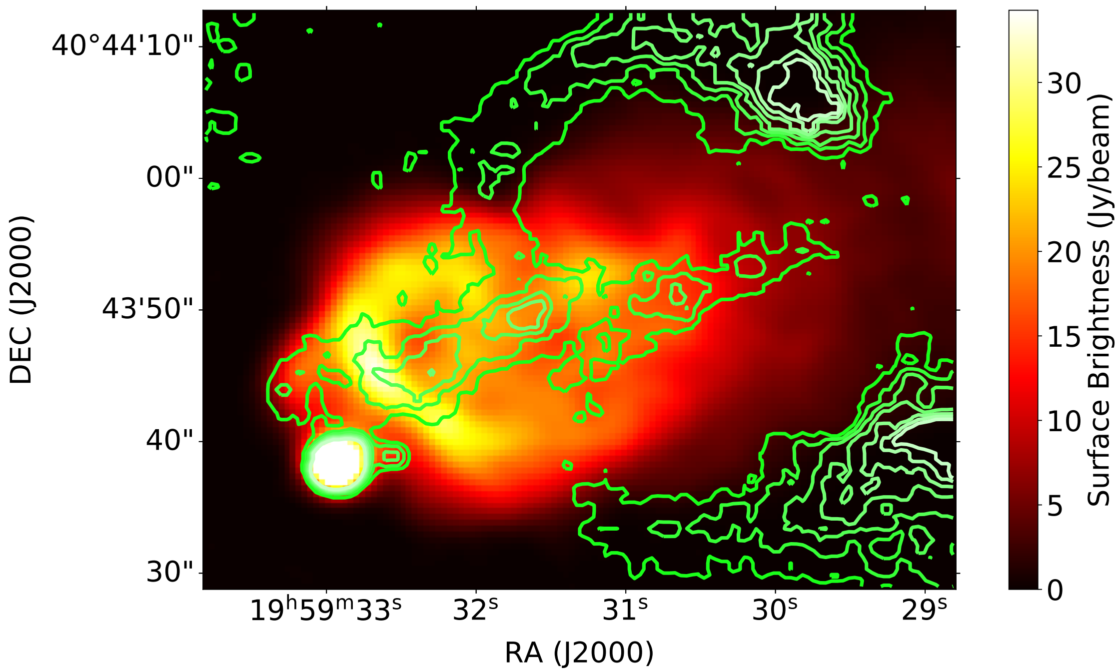

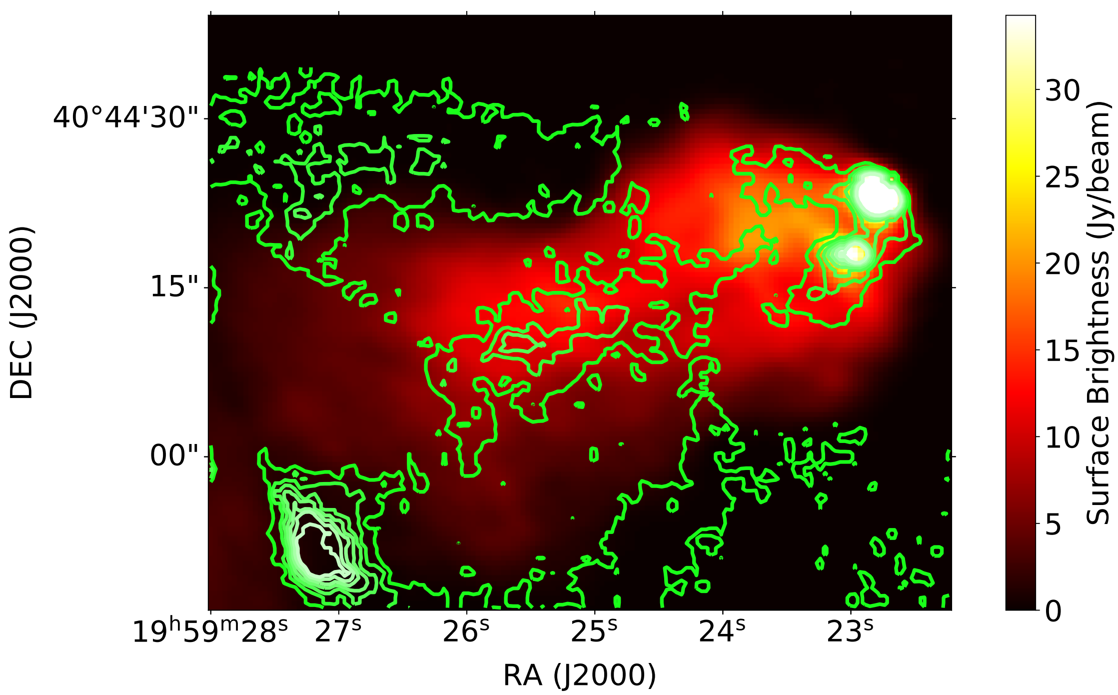

Figure 26: 110-250 MHz LOFAR total intensity with the 0.5-0.7 keV Chandra X-ray contours overlayed of the counter-lobe (upper) and the lobe (lower). The surface brightness in the radio map is saturated at 34 Jy/beam, similar to Figure 19 to show the structure in the lobes more clearly.

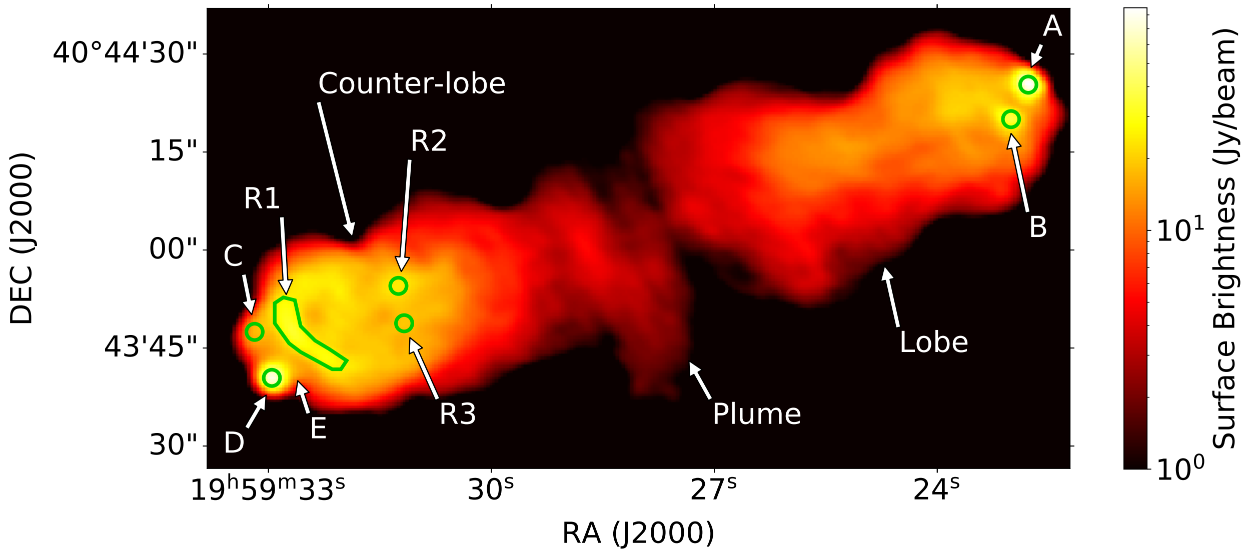

Figure 27: LOFAR radio intensity map, with beams of 2.5 arcsec overlayed to indicate the location of the regions (encircled) where the spectra were taken from. The four known hotspots A, B, D, and E are labeled, as well as the proposed hotspot C. In addition, the extended regions of interest are labeled 'R1', 'R2' and 'R3'.

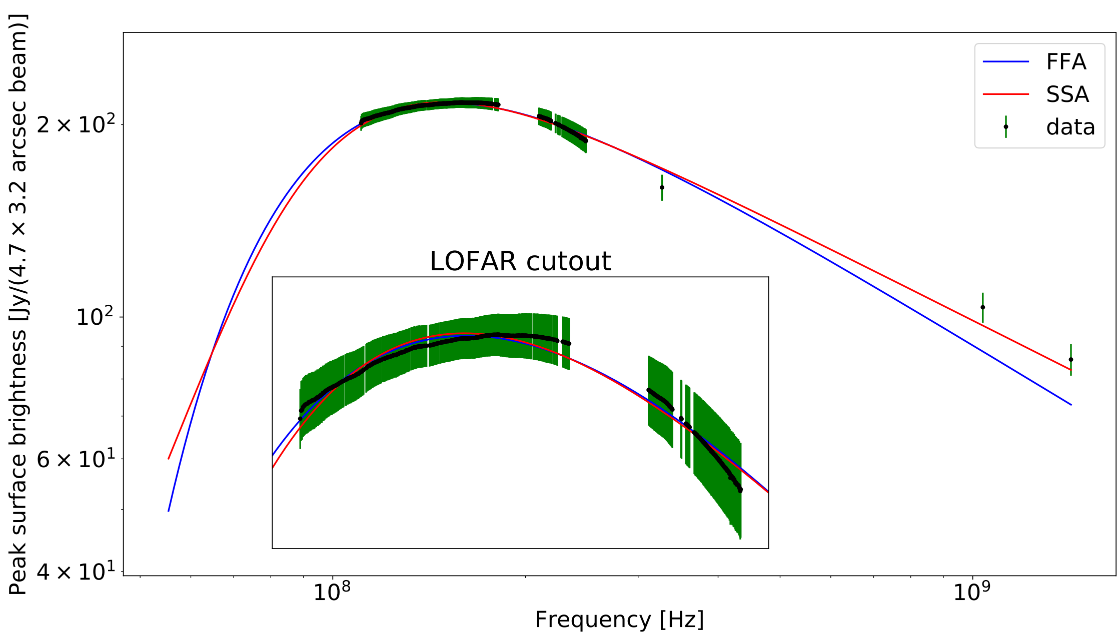

Figure 28: The spectrum of hotspot A between 110 and 1424 MHz, with error bars in green. The surface brightness was measured from the peak pixel value in a circular region around hotspot A with diameter 2.5 arcsec in the images of all separate LOFAR and VLA bands. The best FFA fit is shown in the blue curve, and the best SSA fit in red.

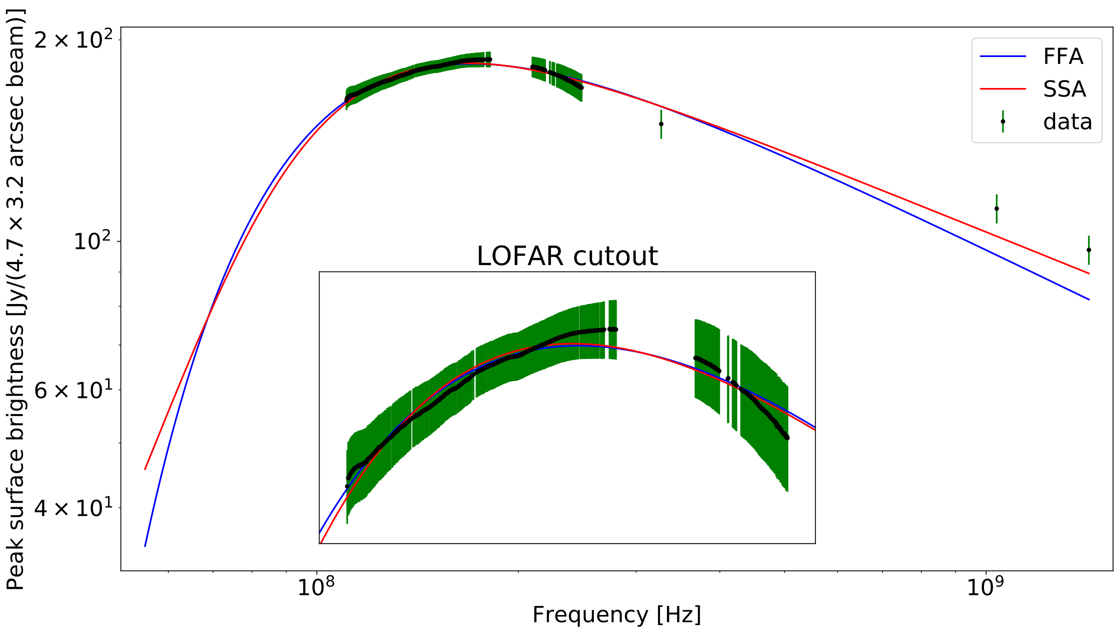

Figure 29: The spectrum of hotspot D between 110 and 1424 MHz, with error bars in green. The surface brightness was measured from the peak pixel value in a circular region around hotspot D with diameter 2.5 arcsec in the images of all separate LOFAR and VLA bands. The best FFA fit is shown in the blue curve, and the best SSA fit in red.

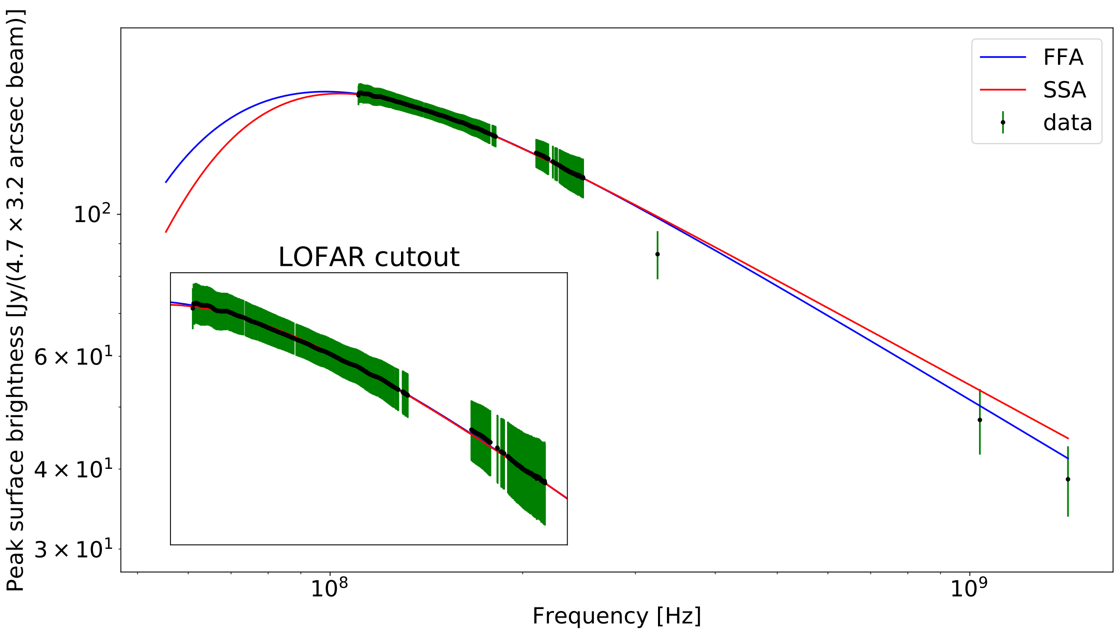

Figure 30: The spectrum of hotspot B between 110 and 1424 MHz, with error bars in green. The surface brightness was measured from the peak pixel value in a circular region around hotspot B with diameter 2.5 arcsec in the images of all separate LOFAR and VLA bands. The best FFA fit is shown in the blue curve, and the best SSA fit in red.

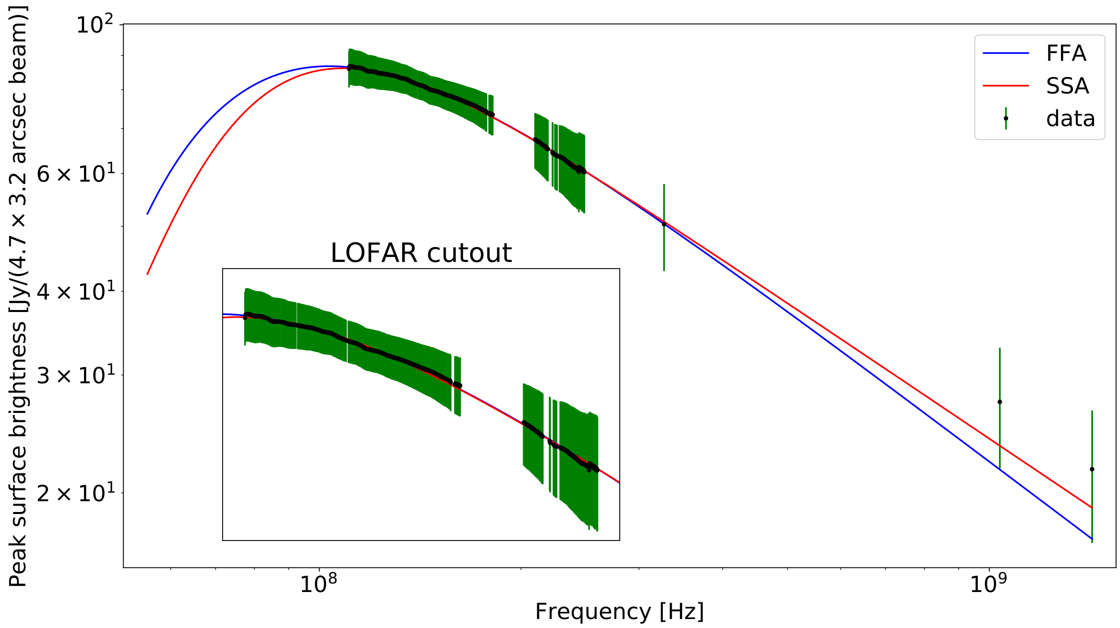

Figure 31: The spectrum of hotspot C between 110 and 1424 MHz, with error bars in green. The surface brightness was measured from the peak pixel value in a circular region around hotspot C with diameter 2.5 arcsec in the images of all separate LOFAR and VLA bands. The best FFA fit is shown in the blue curve, and the best SSA fit in red.

to significant flattening in the observed frequency range, and how the spectra compare to that of generic counter-lobe plasma. The spectra of R1, R2 and R3 can be seen in Figures 32 to 34.

The spectra of hotspot A and D between 110 and 250 MHz in Figures 28 and 29 clearly show how the turnover regimes invert the spectral slopes from the optically thick to the optically thin regime. The spectrum of hotspot A does not seem to have reached a stable slope on the optically thick side of the LOFAR spectrum, as curvature is still seen at near 110 MHz. On the optically thin end it does align fairly well with the 327 MHz data point, which suggests that the slope is more stabilized there. Hotspot D's spectrum shows the same decent alignment with the 327 MHz data point, and on the optically thick end it does seem to be more stabilized than hotspot A, having a spectral index of $\alpha_{\text{thick}} \approx 0.39$ at frequencies below 140 MHz.

The spectra of hotspot B and C shown in Figure 31 show a flattening towards the lowest frequencies of the band. A full turnover can not be seen in either of them, but the spectra near 110 MHz are very close to reaching the peak of a potential turnover. The optically thin spectral index has likely been reached on the high frequency end of the spectrum in both hotspots, but in this case both spectra would be slightly flat compared to that of the generic lobe, as hotspot B is stabilized at $\alpha \approx -0.46$ and hotspot C at $\alpha \approx -0.59$.

From the spectral index and curvature map it was expected that R1 would also have a flat spectrum with significant curvature that could lead to a turnover appearing in the spectrum, but this is not the case. Instead, most of the band seems to be spanned by the optically thin regime, with only slightly curving behaviour, as can be seen in Figure 32. A single power law fit shows that the LOFAR bands have a spectral index of $\alpha = -0.56$, although it must be said that curvature is

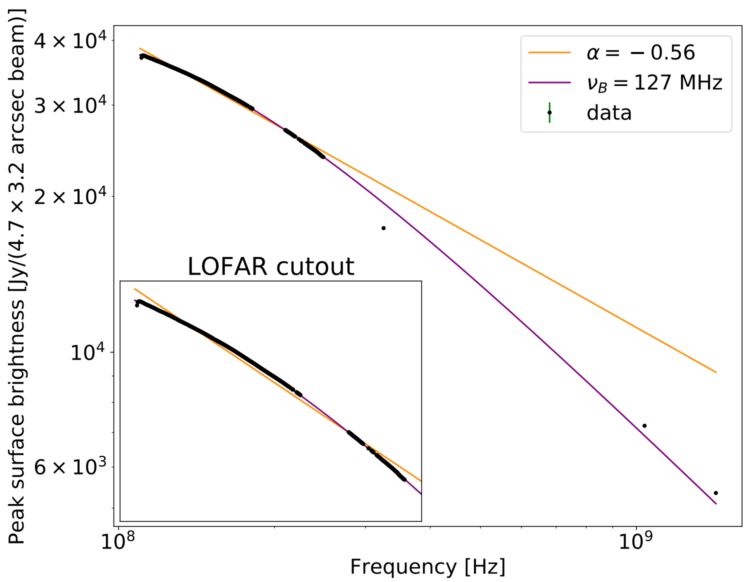

Figure 32: The spectrum of R1 between 110 and 1424 MHz, with error bars in green. The surface brightness was measured summing the pixel values in the demarcated region named R1, in the images of all separate LOFAR and VLA bands. A single power law synchrotron spectrum fit is shown as the orange line, having a spectral index of $\alpha = -0.56$. The best broken power law fit is shown as the purple curve, having break frequency $\nu_{\rm B} = 127\,\rm MHz$.

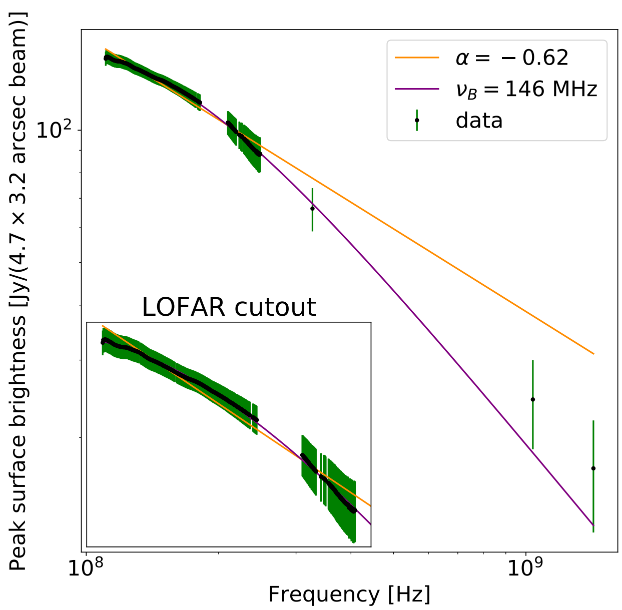

Figure 33: The spectrum of R2 between 110 and 1424 MHz, with error bars in green. The surface brightness was measured from the peak pixel value in a circular region around the filament with diameter 2.5 arcsec in the images of all separate LOFAR and VLA bands. A single power law synchrotron spectrum fit is shown as the orange line, having a spectral index of $\alpha = -0.62$. The best broken power law fit is shown as the purple curve, having break frequency $\nu_{\rm B} = 146\,\rm MHz$.

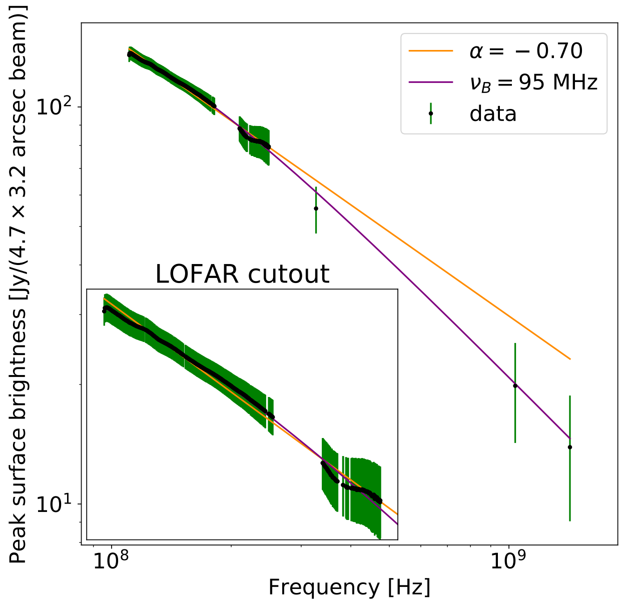

Figure 34: The spectrum of R3 between 110 and 1424 MHz, with error bars in green. The surface brightness was measured from the peak pixel value in a circular region around the filament with diameter 2.5 arcsec in the images of all separate LOFAR and VLA bands. A single power law synchrotron spectrum fit is shown as the orange line, having a spectral index of $\alpha = -0.70$. A broken power law fit is shown as the purple curve, having break frequency $\nu_B = 95$ MHz.

seen data the LOFAR data. A broken power law has also been fitted to account for the curvature, having a break frequency $\nu_B = 127$ MHz. The spectral index below the break is $\alpha_1 = -0.36$, and the spectral index above the break is $\alpha_2 = -0.98$.

Figures 33 and 34 show the spectra of R2 and R3, respectively. As these regions are expected to be of similar ages, and are theorized to be the result of the same shock caused by the moving jet, one might expect their spectra to look very similar. This is seen not to be the case: R2 shows relatively strong curvature and flattening near 110 MHz, but R3 only shows minimal curvature, and is far from reaching a significant flattening at 110 MHz. This curvature causes the LOFAR spectral index of R2 to be flatter than that of R3, being $\alpha = -0.62$ for R2 and $\alpha = -0.70$ for R3. The broken power law fit results in $\nu_B = 146$ MHz, $\alpha_1 = -0.44$ and $\alpha_2 = -1.22$ for R2, while for R3 the result yields $\nu_B = 95$ MHz, $\alpha_1 = -0.52$ and $\alpha_2 = -1.02$. This indicates that the spectrum of R2 is flatter on the low frequency end, but also steeper on the high frequency end than that of R3.

Figures 28 to 31 also show the best fits of both the SSA and FFA model. The SSA model is given by

$$S_\nu = S_{\nu_0}(\nu/\nu_0)^{2.5}\left(1 - \exp\left[-\tau_{\nu_0}(\nu/\nu_0)^{\alpha_{\mathrm{SSA}}-2.5}\right]\right), \tag{32}$$

where S_{ν_0} and τ_{ν_0} are the flux density and optical depth at reference frequency $\nu_0 = 150$ MHz, and α_{SSA} is the optically thin spectral index. It can be seen from the form of the SSA model that the flux density characteristically increases as $\propto \nu^{2.5}$ for low frequencies, and turns over to $\propto \nu^{-\alpha_{\mathrm{SSA}}}$ for high frequencies. The FFA model behaves slightly differently, being given by

$$S_\nu = S_{\nu_0}(\nu/\nu_0)^{\alpha_{\mathrm{FFA}}} \exp\left[-\tau_{\nu_0}(\nu/\nu_0)^{2.1}\right], \tag{33}$$

Table 1: Best fit parameters obtained from the SSA and FFA modelling in the hotspots. All LOFAR and VLA bands were used, hence the fit is dominated by the LOFAR data. S_{ν_0} and τ_{ν_0} are measured at reference frequency ν_0 = 150 MHz.

Hotspot	Model	S_{ν_0} (Jy/beam)	τ_{ν_0}	α_{thin}
A	SSA	726 ± 29	0.35 ± 0.01	-0.51 ± 0.01
	FFA	290 ± 11	0.29 ± 0.01	-0.61 ± 0.01
D	SSA	552 ± 20	0.40 ± 0.01	-0.40 ± 0.01
	FFA	247 ± 2	0.30 ± 0.01	-0.49 ± 0.01
B	SSA	1228 ± 468	0.12 ± 0.01	-0.54 ± 0.01
	FFA	161 ± 1	0.12 ± 0.01	-0.60 ± 0.01
C	SSA	523 ± 33	0.17 ± 0.01	-0.67 ± 0.01
	FFA	94 ± 1	0.17 ± 0.01	-0.76 ± 0.01

where α_{FFA} is the optically thin spectral index for the FFA case (Kassim, 1989). Here it can be seen that the flux density behaves differently at low frequencies than in the SSA model, while the slope of the spectrum will also increase to become steeper than $\propto \nu^{-\alpha_{\text{FFA}}}$ above the turnover, which makes sense from a physical perspective as the effect of FFA is different over the frequency range than that of SSA. The best optimized parameters resulting from the fitting of the SSA and FFA models to the LOFAR + VLA data for all four hotspots can be seen in Table 1.

5 Discussion

5.1 Hotspot A and D

The fit of the SSA models allows for an interpretation of the physical conditions in hotspot A and D. The brightness temperature of a Gaussian source can be calculated from the Rayleigh-Jeans law

$$T_\mathrm{b} = \frac{c^2}{2k\nu^2} I_\nu = \frac{c^2}{2k\nu^2}\frac{S_\nu}{\Delta\Omega}, \tag{34}$$

where k is the Boltzmann constant, and $\Delta\Omega$ is the solid angle of the source, which in the proper units becomes

$$\left[\frac{T_\mathrm{b}}{\mathrm{K}}\right] = 1.22\times10^{12}\left[\frac{S_\nu}{\mathrm{Jy}}\right]\left[\frac{\theta_\mathrm{maj}\theta_\mathrm{min}}{\mathrm{arcsec}^2}\right]^{-1}\left[\frac{\nu}{\mathrm{MHz}}\right]^{-2}, \tag{35}$$

where θ_maj and θ_min are the major and minor axis of the source, respectively. A calculation of the brightness temperature at 150 MHz, for a circular source of 2 kpc in diameter (Wright and Birkinshaw, 2004) gives $T_\mathrm{b} \geq 3.71\times10^9$ K for Hotspot A and $T_\mathrm{b} \geq 3.14\times10^9$ K for hotspot D. These limits are a factor 3 to 4 higher than what was found by McKean et al. (2016), partly due to the smaller beam size of the images. Although this calculation depends on the assumption of a Gaussian source, which will be shown later to probably not be the case, brightness temperatures of this order of magnitude indicate that SSA is likely a significant process in both hotspots.

If the turnover of the spectrum is caused fully by the SSA process, then the peak flux S_p and the peak frequency ν_p of the fit can be used to calculate the magnetic field strength in the region of interest. Assuming that the hotspots are single homogeneous sources in which absorption is dominated by the SSA process, the magnetic field strength is given by

$$\left[\frac{B_\mathrm{SSA}}{\mathrm{Gauss}}\right] \sim \left[\frac{\nu_p}{8\,\mathrm{GHz}}\right]^5\left[\frac{S_\mathrm{p}}{\mathrm{Jy}}\right]^{-2}\left[\frac{\theta_\mathrm{maj}\theta_\mathrm{min}}{\mathrm{milliarcsec}^2}\right]^2(1+z)^{-1}, \tag{36}$$

where z is the redshift of the source (Kellermann, 1981). The resulting magnetic fields are B_SSA = 0.46 G and B_SSA = 1.12 G for Hotspot A and D, respectively. These values are between a factor 2 and 2.5 smaller than the values found by McKean et al. (2016), which were found to be 4 order of magnitude higher than calculated from synchrotron self-Compton and minimum energy fields. Calculations on synchrotron self-Compton emission in X-ray images by Harris et al. (1994) resulted in magnetic fields of 160 and 250 µG for Hotspot A and D, respectively, while minimum energy fields were shown to range from 250 to 350 µG by Carilli et al. (1999).

A couple of nuances and possible explanations can be found for the discrepancy between the calculated magnetic field strengths from the various arguments. One potential issue could be that the calculation has a large dependence on the peak frequency of the synchrotron spectrum. From the full LOFAR + VLA spectra it seems that the peak frequency is relatively well determined, but the zoom on the LOFAR frequencies reveals that a noticeable difference is present between the peak frequency in the LOFAR flux densities and the peak of the SSA fit. However, even an error of 10% in peak frequency would only result in a factor 1.6 difference in the magnetic field strength, and therefore it is likely that there are other effects at play in the forming of this discrepancy.

A potentially bigger issue, already mentioned by McKean et al. (2016), is the assumption of the single homogeneous source, combined with too large a value for the source size. The source size of 2 kpc was chosen because it corresponds roughly to the full MHz regime diameter of hotspots A and D. However, both hotspots are known to contain smaller regions of structure, as 43 GHz maps from Carilli et al. (1999) show that the intensity throughout the hotspots is far from homogeneous, and polarization maps show that also the magnetic field was observed not to be uniform. The 43 GHz images show that Hotspot A and D each consist of 5-7 smaller blobs, seeming to be ordened in two parallel "ridges". The blobs have sizes in the order of 0.2 arcseconds in Hotspot A, and 0.3 arcseconds in Hotspot D. This is an order of magnitude smaller than the resolution of the LOFAR maps, and therefore such structures can not be distinguished in the 110-250 MHz images.

At this point it should be noted that only the Dutch stations were used in this research. The longest baselines of the international stations reach up to 2000 km at the time of writing, which would produce a resolution of $\sim$ 0.17 arcsec at 180 MHz, precise enough to resolve smaller angular scales in the hotspots. Though, calculating the hypothetical size of the regions that would produce the 0.46 and 1.12 G magnetic fields results in sizes of ~ 0.12 arcsec in hotspot A and ~ 0.09 arcsec in hotspot D, assuming that each hotspot contains 6 homogeneous regions of equal size and intensity. It can not be said for sure whether the ILT will be able to resolve the individual regions, as the hotspot structure between 110 and 250 MHz could also be very different from the 43 GHz observations. Resolving the ridge arrangements is potentially feasible, and thus it can be of great value to obtain high quality observations from the international stations.

As said before, FFA could be responsible for the absorption as well. This can either happen inside the hotspots itself, or by a cloud in the line of sight towards the hotspots. In either case, the optical depths found from the FFA fits to the spectra can be used to calculate the electron density in the cloud, as the optical depth is given by

$$\tau_\nu \approx 3.28 \times 10^{-7} \left(\frac{T_e}{10^4\ \mathrm{K}}\right)^{-1.35} \left(\frac{\nu}{\mathrm{GHz}}\right)^{-2.1} \left(\frac{\mathrm{EM}}{\mathrm{pc\ cm^{-6}}}\right), \tag{37}$$

where T_e is the electron temperature, and EM is the emission measure given by

$$\left(\frac{\mathrm{EM}}{\mathrm{pc\ cm^{-6}}}\right) = \int_{\mathrm{los}} \left(\frac{n_e}{\mathrm{cm^{-3}}}\right)^2 \mathrm{d}\left(\frac{s}{\mathrm{pc}}\right) = \langle n_e \rangle^2 \, \Delta\left(\frac{s}{\mathrm{pc}}\right), \tag{38}$$

where the electron density is integrated over the size of the cloud along the line of sight (Mezger and Henderson, 1967). In the case of the absorption happening in our own Milky Way, a cloud size of the order of 1 kpc with an electron temperature of 8000 K can safely be assumed (Kassim, 1989). Using the obtained optical depths from the FFA fit would then result in an average electron density of $n_e = 3.51\,\mathrm{cm^{-3}}$ in the cloud that is in the line of sight towards Hotspot A, and $n_e = 3.55\,\mathrm{cm^{-3}}$ towards Hotspot D. These values are two orders of magnitude larger than the upper limit of $n_e \leq 0.26\,\mathrm{cm^{-3}}$ in the average interstellar medium of the Milky Way found by Kassim (1989). The possibility exists that the line of sight towards Cygnus A is more densely packed with electrons than the average Milky Way ISM. However, these overdensities would then be located in a few specific directions, as other parts of Cygnus A show different amounts of absorption and hence varying optical depths. As a result, the calculated electron density varies a lot between these lines of sight, and this would be too coincidental on these angular scales.

A more likely scenario would be that the hotspot itself contains the absorbing material. In this case, the absorbing material can be assumed to have a size of 2 kpc, and electron densities become $n_e = 2.48\,\mathrm{cm^{-3}}$ and $n_e = 2.51\,\mathrm{cm^{-3}}$ for Hotspot A and D, respectively. These values are nearly four orders of magnitude larger than the upper limit of thermal electrons $n_e < 4 \times 10^{-4}\,\mathrm{cm^{-3}}$ obtained by Dreher et al. (1987) from Faraday rotation measurements. The disclaimer is given that this can be up to two orders of magnitude higher if the magnetic field is strongly anisotropic, which is a possibility as they are likely strongly inhomogeneous as well, but that is still not enough to explain the optical depths observed in the hotspots.

Another solution to the unphysical magnetic fields and electron densities in the hotspots is that multiple processes can be working together in the creation of the turnover. A combination of an absorption process with an LEC can potentially be a working configuration, as the low frequency slopes of the hotspot spectra are flatter than from purely FFA or SSA only. One can imagine that the combined effect of an LEC with SSA could form a spectrum where the LEC induced turnover is at a higher frequency than the SSA dominated regime, leading to a too high peak frequency in the magnetic field calculation. However, the combination of SSA and an LEC will likely not produce physically satisfactory results, as an LEC would reduce the contribution of SSA due to the lack of low energy electrons. A combination of FFA and an LEC is therefore the favoured combination. It was shown by McKean et al. (2016) that an FFA-LEC fit to their data gives a good fit result, although this did not reduce the optical depth enough to reduce the electron densities to acceptable values. Therefore, it remains unclear what process is responsible for the turnover in hotspot A and D.

5.2 Hotspot B and C

The other two spectra that show flattening and a possible turnover are the spectra of Hotspot B and C. Hotspot B shows surface brightnesses similar to hotspot A and D, while hotspot C's surface brightness is relatively low, which is as expected as it is barely identifiable in the radio intensity map. Both spectra flatten significantly near the low end of the observing band.

Following the same approach as in hotspot A and D with a size of 1 kpc, hotspot B has a brightness temperature of $T_{\mathrm{b}} \geq 9.8 \times 10^{9}\,\mathrm{K}$ at 150 MHz, a magnetic field strength of $B_{\mathrm{SSA}} = 9.0 \times 10^{-3}\,\mathrm{G}$ inferred from the SSA fit, and an electron density of $n_e = 2.23\,\mathrm{cm}^{-3}$ from the FFA fit. While the electron density is high, the magnetic field strength takes a more reasonable value, although it is still more than an order of magnitude off. It must be said that its smaller size has a big impact on the outcome of the calculation, but Carilli et al. (1999) also showed that hotspot B contains significantly less substructure. Assuming that it is made up out of 2 blobs of equal intensity, their hypothetical sizes are 0.29 arcsec, which is larger than the resolution of the longest ILT baseline.

While the magnetic field strength in hotspot B seems to take a more reasonable value, this can not be said for hotspot C. Assuming that C is a hotspot of 2 kpc in size, the calculations yield $T_{\mathrm{b}} \geq 1.37 \times 10^{9}\,\mathrm{K}$, $B_{\mathrm{SSA}} = 0.59\,\mathrm{G}$ and $n_e = 1.86\,\mathrm{cm}^{-3}$. This magnetic field strength is still more than three orders of magnitude larger than the 300 μG set for hotspots A and D, which is also likely to be too high for hotspot C even in the case that it is an abandoned hotspot, as expansion and cooling likely have lowered its magnetic field strength. The inferred magnetic field can be lowered to reasonable values if the region consists of 6 substructures of 0.19 arcsec in size each. Finding these substructures and more information on this region in future research can potentially give valuable insight into the nature of Hotspot C, and add a lot to the understanding of the emergence and aging of hotspots within the framework of the dentist drill model.

The shift in the peak frequency in these hotspots compared to hotspots A and D is quite notable. For hotspot C it will be difficult to explain this shift fully from the SSA point of view, as the magnetic field calculations show the same discrepancies as in the hotspots. Also FFA arguments might not be able to fully explain the differences between the spectra, as the factor 2-4 difference in optical depth will only result in an electron density ratio between 1.4 and 2, far less than the multiple orders of magnitude needed to reach reasonable plasma densities. An explanation from the LEC perspective might present a possible solution, as the peak frequency of aging plasma is expected to decrease over time as all electrons radiate away their energy, including the ones just above the LEC. However, from this theory it is expected that the spectra of hotspot B and C would have very different peak frequencies, as hotspot B is an active hotspot and therefore has a low spectral age, while hotspot C is expected to be older. Hence an LEC alone is also not sufficient to fully explain the turnover frequencies of these regions.

To fully confirm that the spectra of these hotspots actually do turn over below 110 MHz, it is necessary to observe their spectrum at lower frequencies. This can be done using LOFAR LBA observations in the 10-80 MHz regime. These were not used for this research for computational reasons, but it can be concluded that they are essential for examining the turnovers that might be hiding in other features than the secondary hotspots. Using the LBA spectra to better contrain the fit parameters, similar approaches can be taken in finding the SSA induced magnetic field and the electron density in case of FFA. There is also the possibility that the LBA spectra do not reveal a turnover, especially in hotspot C. In hotspot B, the absence of a turnover would be puzzling, as it is already known to be a hotspot, and is therefore probably dense enough for absorption processes to happen. It is unlikely that the curvature in these regions is due to aging, as break frequencies in this regime imply ages that are unbelievably high for (relic) hotspots.

For the same reason as for hotspots A and D, data from the ILT stations can potentially be very useful for resolving the underlying structure that might be hiding in hotspot B and C. However, resolving substructure might be more difficult for hotspot C, as the substructures might be too faint. Its faintness might possibly be bypassed by utilizing the spectral index map for the structure identification. This can potentially reveal its structure, even with its lower brightness. This can potentially give a lot of insight in the working of secondary hotspots, which are still poorly understood. The possibility exists that hotspot C is leftover from earlier jet activity. The other explanation would be that hotspot C is a 'splatter spot', in which case it is fed by jet plasma that is being deflected by the magnetic field of the current hotspot (Smith, 1984; Williams and Gull, 1985). However, this would require extreme circumstances given its distance towards hotspot D and E. ILT observations might also reveal a completely different story, as hotspot C could also be interconnected with region R1, in which case the question arises why the flattening in its spectrum is much stronger than that in R1.

5.3 Extended regions

While the spectral index and curvature map clearly showed region R1 spanning across the counterlobe, its spectrum only showed slight curvature. The flattening seen at low frequencies is also not as strong as what is seen in the hotspots. However, the fact that the low-frequency end is flatter, and the high-frequency end is steeper than in the generic lobe plasma, indicates that a shock wave producing event could have created this extended arc of emission. This hypothesis is strengthened by the break frequency of the fitted broken power law: Carilli et al. (1991) found a break frequency in the region of R1 around $\nu_B \sim 20$ GHz using spectral age modelling, implying a spectral age of ~ 1 Myr, assuming a magnetic field of 50 µG. A break frequency of 127 MHz from the fitted power law to the LOFAR + VLA data would imply an age of ~ 13 Myr, which is much higher than the age of the oldest plasma is Cygnus A, even though the broken power law fits the data rather well. Hence, aging as a source of the curvature is ruled out, favouring the particle acceleration theory resulting from the shock.

Compared to the generic plasma, the brightest part of R1 is found to radiate roughly twice as much power per volume unit at the LOFAR frequencies, which also hints to the shock scenario, as a source of energy is required to enhance the brightness in this region. An example of a similar event is seen in radio galaxy 3C 310, which shows a bright ring in one of the lobes that is proposed to be the result of supersonic expanding lobe plasma (Kraft et al., 2012). This shows that shock waves can occur in radio lobes, but more research is required to confirm whether region R1 in Cygnus A has a similar origin as the shock wave in 3C 310.

Region R2 and R3 are located much closer to the nucleus, and are thus expected to be older then R1, as their electron populations have flowed back a larger distance back from the hotspots to the nucleus. Hence, the question arises whether their respective break frequencies of 146 MHz and 95 MHz can be due to aging. Carilli et al. (1991) showed that the break frequency near R2 and R3 lies near $\nu_{\mathrm{B}} \sim 5$ GHz, implying an age of ~ 2 Myr, using the same magnetic field of 50 µG. The break frequencies of 146 MHz and 95 MHz result in a spectral age of $\sim$ 12 Myr for R2 and $\sim$ 15 Myr for R3. Using the same reasoning as for R1, this rules out that the curvature in the LOFAR spectra is caused by synchrotron aging, as the resulting plasma ages are unbelievably high. The hypothesis of a shock caused by the moving jet, as proposed by Perley et al. (1984), is therefore also favoured.

An additional effect that possibly contributes to the curvature in these regions is absorption. Similar to the situation in the hotspots, R1, R2 and R3 can have higher electron densities than the surrounding lobe plasma, as a result of the shock that is propagating through them. This will cause flattening in the spectra, additional to the shock-induced particle acceleration, if this effect is significant in these features. An explanation is then still to be found for differences in curvature and high-frequency spectral index observed between R2 and R3, as these regions are expected to have the same origin, and are thus expected to show comparable behaviour.

Region R3 is seen to be slightly dimmer than R2, following the general trend that the counter-lobe north of the jet is brighter than the south of the jet. Having a single power law slope of $\alpha = -0.70$, which is very common in the counter-lobe, in addition to its small amount of curvature, makes that only the brightness of R3 is an indication that this region is linked to an energy injecting event, like a shock by the moving jet. Similar to R2, ILT observations can potentially give more insight in the properties of region R2 and R3 by obtaining better spatial resolution, which can potentially reveal structural differences between the two regions that are currently not resolved by the Dutch LOFAR stations.

6 Conclusion

The prototype FR-II radio galaxy Cygnus A was imaged using 492 bands of LOFAR HBA observations between 110 and 250 MHz. Using the Dutch stations of the interferometer, an image with an angular resolution of 2.03×1.38 arcsec was produced. The image shows a lot of structure in the synchrotron emitting radio lobes, such as the bright hotspots A, B and D, but also extended features that stretch through large portions of the lobes. Also obtained from the imaging is a spectral index map, showing the change in slope of the spectrum between different regions throughout the source. The spectral index map shows that many of the bright features in the intensity map also have a relatively flat spectral index, even reaching positive slopes in the hotspots. In addition, a spectral curvature map is produced, which shows that the spectra of many of the brightest regions show strong negative curvature, i.e. they steepen towards higher frequencies. This can either be the result of a turnover due to (partial) absorption of the synchrotron emission coming from the considered region, or it can be the result of radiative losses due to synchrotron aging. Absorption is the favoured model, as break frequencies in the LOFAR observing band would imply unbelievably high spectral ages.

A deeper spectral analysis was done on seven regions that stand out in the radio image and the spectral maps. The HBA bands were imaged separately and combined with 3 VLA bands at 327, 1037 and 1424 MHz to be able to retrieve spectra in the seven regions of interest. All the bands were smoothed to the same restoring beam to guarantee that regions experience the same amount of beam smearing in each band, with the assumption that this results in a similar contribution from the lobe in each band. Spectra were taken in secondary hotspots A and D, which have shown in earlier research by McKean et al. (2016) to contain a turnover between 110 and 180 MHz. This turnover is observed again, and curvature due to the turnover is seen still to be present in the 210-250 MHz band. In addition, the spectrum of primary hotspot B was studied, revealing significant flattening towards 110 MHz, but a turnover was not observed directly. The spectrum of a faint blob near the edge of the counter-lobe was also investigated, as it shows hotspot-like behaviour in the spectral maps, making it a potential relic hotspot candidate. In this research the region was given the name hotspot C, with the disclaimer that more research needs to be carried out to confirm whether this is in fact a hotspot, as both the splatter spot and dentist drill model have difficulties explaining its location. As with hotspot B, the spectrum of hotspot C did not show a direct turnover, but its curvature is nearly flat near 110 MHz, indicating that a turnover can possibly be observed at lower frequencies. The remaining three spectra were taken from three extended features in the counter-lobe, which are thought to be associated with particle acceleration through shocks.

In the circular shaped regions (hotspots A, D, B & C), the fitting of a synchrotron self-absorption model allowed for the calculation of the magnetic field strength, and from a free-free absorption model fit an electron density in the free-free absorbing cloud be estimated. Magnetic field strengths inferred from the fitted SSA models in hotspot A and D are found to be 0.46 G and 1.12 G, repectively, which are more than three orders of magnitudes higher than values obtained from equipartition and Faraday rotation arguments (Harris et al., 1994; Carilli et al., 1999). This discrepancy was also found by McKean et al. (2016), and the proposed solution is that hotspot A and D are not homogeneous, but instead consist of smaller highly magnetic blobs, as observed by Carilli et al. (1999). Six blobs of sizes ≤ 0.09 arcsec in both hotspots are needed to obtain reasonable magnetic field strengths in the individual blobs.

Using the same calculation for hotspot B, a magnetic field strength of 9×10^{-3} G was inferred. This calculation takes into account that hotspot B has a smaller angular size, which has a big impact on the numerical outcome of the calculation. Although this magnetic field seems more feasible than those of hotspot A and D, it is still more than an order of magnitude above the expected strength from (Harris et al., 1994; Carilli et al., 1999). Hotspot B is also known to contain substructure that can possibly solve the discrepancy, but this can be more difficult because of its smaller size. The magnetic field in hotspot C is found to be 0.59 G, which of the same order of magnitude as that of hotspot A and D.

The fitted FFA models resulted in an electron density around $n_e \sim 3.5\,\mathrm{cm}^{-3}$ for both hotspot A and D in the situation where the absorbing medium is located outside the Cygnus cluster, and around $n_e \sim 2.5\,\mathrm{cm}^{-3}$ if the hotspots themselves are responsible for the FFA. For hotspots B and C these values were found to be slightly lower, yielding $\sim 2\,\mathrm{cm}^{-3}$. All these values are multiple orders of magnitude above numbers estimated from Kassim (1989) and Dreher et al. (1987). A combination of SSA and FFA models, and the addition of an LEC might provide a more satisfactory explanation why the slope of the spectrum of hotspot A and D below the turnover seems to be stabilized slightly flatter than both the SSA and FFA model, as an LEC flattens the low-frequency spectrum.

The brightness and the size of the bright arc in the counter-lobe (R1) raise questions on its origin and properties. Its spectrum is seen to flatten slightly, but not as extreme as what is expected from the spectral curvature map. Its spectral index is signifcantly flatter than that of generic counter-lobe plasma, suggesting that energy has been fed into the region by some process. A shock wave is a candidate, as this is known to be capable of accelerating the electrons, boosting their kinetic energy. Radio galaxy 3C 310 shows a similar event in one of its lobes, indicating that such events could be rather common in radio galaxies. Two additional filamentary structures oriented parallel to the counter-jet (R2 and R3) have been studied to see if any conclusions can be drawn on their physical origin. From their location in the counter-lobe and their orientation towards the jet, their spectra were expected to be comparable, with only a difference in intensity. However, R2 (north) shows significantly more curvature in the LOFAR observing range than R3 (south), and its optically thin spectral index is significantly steeper. Future observations can potentially reveal the physical nature of these features, and reveal structure differences that can explain the differences in their spectra.

6.1 Future recommendations

Observations by the International LOFAR Telescope can be of great value in studying the origin of hotspots A, D, B and C. ILT baselines can reach angular resolutions of the order 0.2 arcseconds, which might not be able to resolve the individual blobs contained in hotspot A and D fully, but it will likely reveal the arrangement of the blobs for a comparison with GHz images. Observing the structure that is potentially hiding in the hotspots at LOFAR observing frequencies can have implications for the inferred physical parameters that result from SSA and FFA modelling, and thereby add to our understanding of hotspots. Obtaining spectra of the individual blobs will be very ambitious given the resolution of the longest ILT baselines.

Hotspot B has a smaller angular size, and is likely to contain fewer pieces of substructure. This is not necessarily problematic, as the hypothetical blobs sizes have sizes that can be resolved by ILT observations. Resolving these substructures with can be of great value to obtain a better understanding of the internal working of hotspot B, and thereby to our understanding of the working of hotspots in general. Resolving similar substructures in hotspot C can help in the search of its physical origin, as observing similar arrangements of smaller blobs can help finding out whether this is a relic hotspot. The ILT observations also have the potential of giving a more detailed look into the extended structures R1, R2 and R3. These can help constrain the physical properties of these regions, and thereby add to the understanding of the events that produce these structures in radio lobes, such as shocks.

LOFAR LBA data can also be very helpful to complete the picture in the investigated regions. It was shown that the regime below the turnovers of hotspot A and D is not fit ideally by the absorption models, and observing the behaviour in the 10-80 MHz regime can provide better constraints on the parameters of the fitted models, and thereby the physical conditions in hotspots A and D. In the case of hotspot B and C, LBA observations can likely reveal turnovers that might exist below 110 MHz. Similar to hotspots A and D, this can provide better constraints on the model parameters in hotspot B. In hotspot C this could be a clue, although not a confirmation, in the question whether it is a relic hotspot resulting from earlier jet activity. Similarly, R1, R2 and R3 might also show more absorption features at these wavelengths, constraining the physical conditions in these features.

References

Kazunori Akiyama, Antxon Alberdi, Walter Alef, Keiichi Asada, Rebecca Azulay, Anne-Kathrin Baczko, David Ball, Mislav Baloković, John Barrett, Dan Bintley, et al. First m87 event horizon telescope results. iv. imaging the central supermassive black hole. *The Astrophysical Journal Letters*, 875(1):L4, 2019.

Robert Antonucci. Unified models for active galactic nuclei and quasars. *Annual review of astronomy and astrophysics*, 31:473–521, 1993.

Arwa Dabbech. VLA L-band images. NRAO FTP server, February 2021. Accessed on February 11, 2021.

Walter Baade and Rudolph Minkowski. Identification of the radio sources in cassiopeia, cygnus a, and puppis a. In *Classics in Radio Astronomy*, pages 251–272. Springer, 1954.

Peter D Barthel. Is every quasar beamed? *The Astrophysical Journal*, 336:606–611, 1989.

RD Blandford and Arieh Königl. Relativistic jets as compact radio sources. *The Astrophysical Journal*, 232:34–48, 1979.

RD Blandford and MJ Rees. A 'twin-exhaust'model for double radio sources. *Monthly Notices of the Royal Astronomical Society*, 169(3):395–415, 1974.

Roger D Blandford and Roman L Znajek. Electromagnetic extraction of energy from kerr black holes. *Monthly Notices of the Royal Astronomical Society*, 179(3):433–456, 1977.

Katherine M Blundell and Steve Rawlings. Spectral ageing: a new age perspective. *arXiv preprint astro-ph/0209372*, 2002.

CL Carilli, RA Perley, JW Dreher, and JP Leahy. Multifrequency radio observations of cygnus a-spectral aging in powerful radio galaxies. *The Astrophysical Journal*, 383:554–573, 1991.

CL Carilli, JD Kurk, Paul P van der Werf, RA Perley, and GK Miley. High-resolution millimeter and infrared observations of the hot spots of cygnus a. *The Astronomical Journal*, 118(6):2581, 1999.

Bradley W Carroll and Dale A Ostlie. *An introduction to modern astrophysics*. Cambridge University Press, 2017.

Gayoung Chon, Hans Böhringer, Martin Krause, and Joachim Trümper. Discovery of an x-ray cavity near the radio lobes of cygnus a indicating previous agn activity. *Astronomy & Astrophysics*, 545:L3, 2012.

Tim J Cornwell. Multiscale clean deconvolution of radio synthesis images. *IEEE Journal of selected topics in signal processing*, 2(5):793–801, 2008.

TJ Cornwell. A method of stabilizing the clean algorithm. *Astronomy and Astrophysics*, 121: 281–285, 1983.

Heber Doust Curtis. Descriptions of 762 nebulae and clusters photographed with the crossley reflector. *Publications of Lick Observatory*, 13:9–42, 1918.

JW Dreher, CL Carilli, and RA Perley. The faraday rotation of cygnus a-magnetic fields in cluster gas. *The Astrophysical Journal*, 316:611–625, 1987.

DO Edge, JR Shakeshaft, WB McAdam, JE Baldwin, and S Archer. A survey of radio sources at a frequency of 159 mc/s. *MmRAS*, 68:37–60, 1959.

RD Ekers, R Fanti, C Lari, and P Parma. Ngc326—a radio galaxy with a precessing beam? *Nature*, 276(5688):588–590, 1978.

Bernard L Fanaroff and Julia M Riley. The morphology of extragalactic radio sources of high and low luminosity. *Monthly Notices of the Royal Astronomical Society*, 167(1):31P–36P, 1974.

Edward Arthur Fath. The spectra of some spiral nebulae and globular star clusters. *Publications of the Astronomical Society of the Pacific*, 21(126):138–143, 1909.

LEH Godfrey, GV Bicknell, JEJ Lovell, David L Jauncey, J Gelbord, DA Schwartz, HL Marshall, M Birkinshaw, M Georganopoulos, DW Murphy, et al. A multiwavelength study of the high surface brightness hot spot in pks 1421–490. *The Astrophysical Journal*, 695(1):707, 2009.

JP Hamaker, JD Bregman, and RJ Sault. Understanding radio polarimetry. i. mathematical foundations. *Astronomy and Astrophysics Supplement Series*, 117(1):137–147, 1996.

PJ Hargrave and M Ryle. Observations of cygnus a with the 5-km radio telescope. *Monthly Notices of the Royal Astronomical Society*, 166(2):305–327, 1974.

DE Harris, CL Carilli, and RA Perley. X-ray emission from the radio hotspots of cygnus a. *Nature*, 367(6465):713–716, 1994.

JA Högbom. Aperture synthesis with a non-regular distribution of interferometer baselines. *Astronomy and Astrophysics Supplement Series*, 15:417, 1974.

JA Hogbom and It Carlsson. Observations of the structure and polarization of intense extragalactic radio sources at 1415 mhz. *Astronomy and Astrophysics*, 34:341–354, 1974.

Edwin Hubble. A relation between distance and radial velocity among extra-galactic nebulae. *Proceedings of the national academy of sciences*, 15(3):168–173, 1929.

RC Jennison and MK Das Gupta. Fine structure of the extra-terrestrial radio source cygnus i. *Nature*, 172(4387):996–997, 1953.

Nikolai S Kardashev. Nonstationarity of spectra of young sources of nonthermal radio emission. *Soviet Astronomy*, 6:317, 1962.

Namir Kassim, RA Perley, CL Carilli, DE Harris, and WC Erickson. Low frequency observations of cygnus a. *Cygnus A–Studay of a Radio Galaxy*, page 182, 1996.

Namir E Kassim. Low-frequency observations of galactic supernova remnants and the distribution of low-density ionized gas in the interstellar medium. *The Astrophysical Journal*, 347:915–924, 1989.

Kenneth I Kellermann. The spectra of non-thermal radio sources. *The Astrophysical Journal*, 140: 969, 1964.

KL Kellermann. Compact radio sources. *Ann. Rev. Astron. Astrophys*, 373:410, 1981.

Andrew King and Ken Pounds. Powerful outflows and feedback from active galactic nuclei. *Annual Review of Astronomy and Astrophysics*, 53:115–154, 2015.

Uli Klein. Dwarf galaxies and the magnetisation of the igm. *Proceedings of the International Astronomical Union*, 4(S255):167–172, 2008.

RP Kraft, M Birkinshaw, PEJ Nulsen, DM Worrall, JH Croston, WR Forman, MJ Hardcastle, C Jones, and SS Murray. An active galactic nucleus driven shock in the intracluster medium around the radio galaxy 3c 310. *The Astrophysical Journal*, 749(1):19, 2012.

T Joseph W Lazio, Aaron S Cohen, Namir E Kassim, RA Perley, WC Erickson, CL Carilli, and PC Crane. Cygnus a: A long-wavelength resolution of the hot spots. *The Astrophysical Journal Letters*, 642(1):L33, 2006.

JP Leahy and AG Williams. The bridges of classical double radio sources. *Monthly Notices of the Royal Astronomical Society*, 210(4):929–951, 1984.

JP Leahy, TWB Muxlow, and PW Stephens. 151-mhz and 1.5-ghz observations of bridges in powerful extragalactic radio sources. *Monthly Notices of the Royal Astronomical Society*, 239 (2):401–440, 1989.

Malcolm S Longair. *High energy astrophysics*. Cambridge university press, 2010.

JP McKean, LEH Godfrey, S Vegetti, Michael W Wise, R Morganti, MJ Hardcastle, D Rafferty, J Anderson, IM Avruch, R Beck, et al. Lofar imaging of cygnus a–direct detection of a turnover in the hotspot radio spectra. *Monthly Notices of the Royal Astronomical Society*, 463(3):3143–3150, 2016.

BR McNamara and PEJ Nulsen. Heating hot atmospheres with active galactic nuclei. *Annu. Rev. Astron. Astrophys.*, 45:117–175, 2007.

BR McNamara and PEJ Nulsen. Mechanical feedback from active galactic nuclei in galaxies, groups and clusters. *New Journal of Physics*, 14(5):055023, 2012.

PG Mezger and AP Henderson. Galactic h ii regions. i. observations of their continuum radiation at the frequency 5 ghz. *The Astrophysical Journal*, 147:471, 1967.

Albert A Michelson and Francis G Pease. Measurement of the diameter of alpha-orionis by the interferometer. *Proceedings of the National Academy of Sciences of the United States of America*, 7(5):143, 1921.

P Parma, M Murgia, R Morganti, A Capetti, HR De Ruiter, and R Fanti. Radiative ages in a representative sample of low luminosity radio galaxies. *arXiv preprint astro-ph/9812413*, 1998.

Daniel A Perley, Richard A Perley, Vivek Dhawan, and Christopher L Carilli. Discovery of a luminous radio transient 460 pc from the central supermassive black hole in cygnus a. *The Astrophysical Journal*, 841(2):117, 2017.

RA Perley, JW Dreher, and JJ Cowan. The jet and filaments in cygnus a. *The Astrophysical Journal*, 285:L35–L38, 1984.

S Pyrzas, KC Steenbrugge, and KM Blundell. Multiwavelength study of cygnus a-v. the hotspots in the lobe. *Astronomy & Astrophysics*, 574:A30, 2015.

George B Rybicki and Alan P Lightman. *Radiative processes in astrophysics*. John Wiley & Sons, 2008.

EE Salpeter. Accretion of interstellar matter by massive objects. *Publications*, 1:165, 1964.

PAG Scheuer. Morphology and power of radio sources. In *Symposium-International Astronomical Union*, volume 97, pages 163–165. Cambridge University Press, 1982.

PAG Scheuer and PJS Williams. Radio spectra. *Annual Review of Astronomy and Astrophysics*, 6:321, 1968.

Maarten Schmidt. 3c 273: a star-like object with large red-shift. *Nature*, 197(4872):1040–1040, 1963.

M Lerato Sebokolodi, Rick Perley, Jean Eilek, Chris Carilli, Oleg Smirnov, Robert Laing, Eric Greisen, and Michael Wise. A wideband polarization study of cygnus a with the jvla. i: The observations and data. *arXiv preprint arXiv:2009.06554*, 2020.

Carl K Seyfert. Nuclear emission in spiral nebulae. *The Astrophysical Journal*, 97:28, 1943.

O Shemmer, H Netzer, R Maiolino, E Oliva, S Croom, E Corbett, and L Di Fabrizio. Near-infrared spectroscopy of high-redshift active galactic nuclei. i. a metallicity-accretion rate relationship. *The Astrophysical Journal*, 614(2):547, 2004.

Joseph Silk and Gary A Mamon. The current status of galaxy formation. *Research in Astronomy and Astrophysics*, 12(8):917, 2012.

MD Smith. Jet deflection and hairpin bends. *Monthly Notices of the Royal Astronomical Society*, 211(4):767–773, 1984.

Bradford Snios, Paul EJ Nulsen, Michael W Wise, Martijn De Vries, Mark Birkinshaw, Diana M Worrall, Ryan T Duffy, Ralph P Kraft, Brian R McNamara, Chris Carilli, et al. The cocoon shocks of cygnus a: Pressures and their implications for the jets and lobes. *The Astrophysical Journal*, 855(1):71, 2018.

Katrien C Steenbrugge and Katherine M Blundell. Multiwavelength study of cygnus a–i. precession and slow jet speeds from radio observations. *Monthly Notices of the Royal Astronomical Society*, 388(4):1457–1464, 2008.

Katrien C Steenbrugge, Katherine M Blundell, and Peter Duffy. Multiwavelength study of cygnus a–ii. x-ray inverse-compton emission from a relic counterjet and implications for jet duty cycles. *Monthly Notices of the Royal Astronomical Society*, 388(4):1465–1472, 2008.

Katrien C Steenbrugge, Ian Heywood, and Katherine M Blundell. Multiwavelength study of cygnus a–iii. evidence for relic lobe plasma. *Monthly Notices of the Royal Astronomical Society*, 401(1): 67–76, 2010.

KC Steenbrugge, KM Blundell, and S Pyrzas. Multiwavelength study of cygnus a-iv. proper motion and location of the nucleus. *Astronomy & Astrophysics*, 563:A131, 2014.

Anthony Richard Thompson, James M Moran, George Warner Swenson, et al. *Interferometry and synthesis in radio astronomy*. Wiley New York et al., 1986.

C Megan Urry and Paolo Padovani. Unified schemes for radio-loud active galactic nuclei. *Publications of the Astronomical Society of the Pacific*, 107(715):803, 1995.

Michael P van Haarlem, Michael W Wise, AW Gunst, George Heald, John P McKean, Jason WT Hessels, A Ger de Bruyn, Ronald Nijboer, John Swinbank, Richard Fallows, et al. Lofar: The low-frequency array. *Astronomy & astrophysics*, 556:A2, 2013.

KL Visnovsky, Christopher D Impey, CB Foltz, PC Hewett, RJ Weymann, and SL Morris. Radio properties of optically selected quasars. *The Astrophysical Journal*, 391:560–568, 1992.

AG Williams and SF Gull. Multiple hotspots in extragalactic radio sources. *Nature*, 313(5997): 34–36, 1985.

MCH Wright and M Birkinshaw. The hot spots of cygnus a at 230 ghz. *The Astrophysical Journal*, 614(1):115, 2004.

Yakov Borisovich Zel'dovich. The fate of a star, and the liberation of gravitation energy in accretion. In *Doklady Akademii Nauk*, volume 155, pages 67–69. Russian Academy of Sciences, 1964.

Frederik Zernike. The concept of degree of coherence and its application to optical problems. *Physica*, 5(8):785–795, 1938.