



Universiteit
Leiden

The Netherlands

Cosmic depth and detail: advancing LOFAR imaging workflows to unveil the deep high-resolution universe

Jong, J.M.G.H.J. de

Citation

Jong, J. M. G. H. J. de. (2025, May 9). *Cosmic depth and detail: advancing LOFAR imaging workflows to unveil the deep high-resolution universe*. Retrieved from <https://hdl.handle.net/1887/4245860>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4245860>

Note: To cite this publication please use the final published version (if applicable).

Introduction

1.1. Birth of radio astronomy

For most of our history, visible light was the only form of radiation that we were aware of, not realizing it was just a small part of a vast electromagnetic spectrum. This changed when infrared radiation was discovered in the early 1800s (Herschel, 1800; Rowan-Robinson, 2013). Over time, new regions of the electromagnetic spectrum were uncovered, providing valuable insights not only into terrestrial phenomena but also into extraterrestrial processes. Among these newly explored regions was radio, which turned out to open a new window into understanding the Universe.

As with many discoveries, the first detections of radio waves of extraterrestrial origin were made by ‘accident’. Early reports of mysterious ‘hissing’ noises at radio frequencies were documented in the 1920s and 1930s (e.g. Oswald, 1930; Arakawa, 1936), which were later identified as being caused by solar activity (e.g. Heightman, 1936; Ham, 1975; Kellermann et al., 2020). The first confirmed detection of radio emission from an extraterrestrial source came in 1933, when physicist and engineer Karl Jansky, while investigating interference in telecommunications, identified signals originating from the centre of the Milky Way at a frequency of 20.5 MHz (Jansky, 1933). Inspired by this discovery, Grote Reber – an engineer, ‘amateur’ radio astronomer, and radio operator – confirmed Jansky’s source of radio emission at 162 MHz. He also detected radio signals from the Andromeda Galaxy, as well as from other areas in the sky, including the Cygnus and Cassiopeia regions (Reber, 1940, 1944). The advancements in radar technology during the Second World War inspired the development of radio astronomy in the postwar era (e.g. Sullivan, 2009; Elbers, 2015; Kellermann et al., 2020). It was therefore after the war that interest in radio astronomy grew, with more engineers and astronomers dedicating their efforts

to capturing and studying signals from the radio sky.

1.2. Radio emission across the Universe

The next most notable discoveries in the radio were the brightest radio sources in the sky, which are grouped into the so-called ‘A-team’. This group includes radio galaxies such as Cygnus A, Virgo A, Hercules A, and Centaurus A, as well as supernova remnants Taurus A, Cassiopeia A, and Puppis A (e.g. Hey et al., 1946; Bolton & Stanley, 1948b,a; Mills, 1952b,a; Baade & Minkowski, 1954). These detections sparked the interest in conducting large-scale surveys, beginning with the Cambridge catalogues (Shakeshaft et al., 1955; Edge et al., 1959; Pilkington & Scott, 1965), and eventually leading to large surveys utilizing more advanced instruments, such as the Faint Images of the Radio Sky at Twenty centimetres (FIRST; Becker et al., 1995), the Westerbork Northern Sky Survey (WENSS; Rengelink et al., 1997), the NRAO VLA Sky Survey (NVSS; Condon et al., 1998), the Galactic and Extra-Galactic All-Sky MWA Survey (GLEAM; Wayth et al., 2015) the TIFR GMRT Sky Survey (TGSS; Intema et al., 2017), and the LOFAR Two-metre Sky Survey (LoTSS; Shimwell et al., 2017).

Over time, various discoveries were made at radio frequencies as well, including discoveries of the 21-cm line (Ewen & Purcell, 1951), quasars (Hoyle & Fowler, 1963; Schmidt, 1963), the cosmic microwave background (CMB; Penzias & Wilson, 1965), rapidly rotating neutron stars (pulsars; Hewish et al., 1968), radio haloes (Willson, 1970), fast radio bursts (FRBs; Lorimer et al., 2007), brown dwarfs (Berger et al., 2001), and more. While each of these objects deserves its own detailed description, we highlight in the following subsections only the background of objects most relevant to the work presented in this thesis.

1.2.1. Radio-loud active galactic nuclei

Among the brightest compact sources that were discovered in the early 1960s were a class of sources called ‘quasi-stellar’ (e.g. Hoyle & Fowler, 1963), which became later known as ‘quasars’ (e.g. Schmidt, 1970). After linking these objects to optical counterparts with known redshifts (e.g. Matthews & Sandage, 1963), it turned out that these sources corresponded to distant galaxies, rather than being from stellar origin (Schmidt, 1963, 1968). Eventually, the widely accepted explanation for the origin of these objects converged to one in which large amounts of material get pulled towards the black hole at the core of galaxies, forming a rotating hot accretion disk around it, from where vast amounts of energy are emitted across the electromagnetic spectrum (e.g. Lynden-Bell, 1969; Netzer, 2013; Padovani et al., 2017). These regions are now known as active galactic nuclei (AGN). The primary mechanism

behind the radio emission from galaxies with AGN is charged particles spiralling at relativistic speeds through magnetic fields, generating non-thermal synchrotron radiation.

When AGN emit a significant portion of their energy through radio waves they are called radio-loud AGN (RLAGN; e.g. Peacock et al., 1986). These types of objects are typically associated with jet-like structures originating from their cores (e.g. Padovani et al., 2017). These jets are energetic outflows extending from kpc up to Mpc scales, far from their host galaxies (e.g. Northover, 1973; Longair et al., 1973; Willis et al., 1974; Bridle & Perley, 1984; Oei et al., 2022, 2024a), and inject vast amounts of energy into the intergalactic medium (IGM) and intra-cluster medium (ICM). This affects galaxy formation within the jet's environment, since it can, for instance, prevent gas from cooling and collapsing to form new stars (e.g. Croton et al., 2006; Nesvadba et al., 2008; Fabian, 2012; McNamara & Nulsen, 2012; Dubois et al., 2013; Tremblay et al., 2016; Morganti, 2017).




Figure 1.1: Example of an FRI (*left panel*) and FRII (*right panel*), indicating the host galaxy and their jets. Images are taken from Hardcastle & Croston (2020).

Based on their morphologies, RLAGN with jets are separated into two different main morphological classes: FRI and FRII radio galaxies (Fanaroff & Riley, 1974). The FRI type has the brightest region near the centre and gradually fades out towards the edge (core-bright), while the FRII type has its brightest hotspots at the edges of the jets (edge-bright). An example of an FRI and FRII are shown in Figure 1.1. It is currently thought that the distinction between these two types of sources stems from the power of their radio jets in combination with the environments they inhabit (e.g. Bicknell, 1995; Kaiser & Best, 2007; Mingo et al., 2019). The FRIs correspond to jets that fail to push through the dense environments in which they

live, causing the jets to diffuse as they extend outward. In contrast, FRIIs have jets powerful enough to sustain their collimated jets while pushing through their environment. As a result, radio galaxies with FRI-type jets are typically found in denser environments like galaxy clusters, while galaxies with FRII-type jets tend to reside often in more isolated regions with abundant gas, providing the resources necessary for their jets to remain undisrupted (e.g. Bicknell, 1994; Croston et al., 2019).

The physical origin of the binary division of RLAGN into FRIs and FRIIs has frequently been challenged due to selection biases introduced by the sensitivity and resolution limits of radio telescopes (e.g. Best, 2009; Singal & Rajpurohit, 2014; Mingo et al., 2019; Magliocchetti, 2022; Mostert et al., 2024), as well as the identification of new populations beyond the binary FR-classification. These include FR0s, which are sources that are about 30 times more core-dominated than ‘standard’ FRIs due to a lack of extended emission (Baldi et al., 2015). There are also hybrid morphologies, which are a mix of both FRI and FRII morphologies (Gopal-Krishna & Wiita, 2000). These ongoing revisions and new classes challenge our understanding of radio galaxy jets, their physical origins, and evolution, leaving us with questions that remain partially or entirely unanswered. Most recent studies suggest that there is no link between different types of accretion and a radio galaxy having FRI or FRII type of jets (Mingo et al., 2022). However, could there be other mechanisms within the host galaxy that may still contribute to jet disruption? Or is jet disruption only driven by changes in the larger-scale environment of the galaxy? Do the FRI and FRII populations evolve differently as a function of redshift and radio power? How does this connect to our understanding of the evolution of galaxies or, on an even larger scale, the cosmic web? These questions highlight the need for multi-wavelength data, including large samples of FRI and FRII radio galaxies from deep radio observations, using different resolutions to strike a balance between surface brightness sensitivity and the ability to resolve their jets. We touch upon this in Chapter 3.

1.2.2. (Pre-)merging galaxy clusters

FRIs and FRIIs with sizes ≥ 0.7 Mpc (e.g. Dabhade et al., 2020)¹ are typically classified as giant radio(-loud) galaxies (GRGs). These objects can serve as tracers of the large-scale structure of the cosmic web in our Universe (e.g. Malarecki et al., 2015; Oei et al., 2024b). At even larger scales, galaxy clusters are another important tracer of the large-scale structure. These clusters of galaxies can span several Mpc and contain masses up to $10^{15}M_{\odot}$, making them the largest and most massive

¹Or in other studies with size thresholds ≥ 1 Mpc (e.g. Lara et al., 2001).

gravitationally bound structures in the Universe. It is therefore that the formation, growth, and merging of galaxy clusters, play an important role in the evolution of the large-scale structure of our Universe (e.g. Springel et al., 2005; Kravtsov & Borgani, 2012).

The mergers of galaxy clusters are very energetic events, releasing $\sim 10^{65}$ ergs (e.g. Markevitch & Vikhlinin, 2007). When the ICMs of merging clusters collide, the gas is compressed and heated through shocks and adiabatic compression, raising the temperature to levels that produce X-ray emission (e.g. Böhringer & Werner, 2010). During mergers, shocks and turbulence accelerate cosmic-ray electrons to relativistic speeds, generating non-thermal synchrotron emission through interactions with magnetic fields, which is most prominently detected at low radio frequencies (e.g. van Weeren et al., 2021).

The two main types of extended radio emission associated with galaxy clusters have been classified as radio (mini-)haloes and relics (e.g. Willson, 1970; Jaffe & Rudnick, 1979; Burns et al., 1992; Brunetti et al., 2009; Cassano et al., 2010; van Weeren et al., 2013, 2017; Bonafede et al., 2022). Radio haloes are often located at the centres of galaxy clusters and span sizes on the order of Mpc. Their origin is currently explained by two main models: turbulent reacceleration of cosmic-ray electrons within the ICM (De Young, 1992; Brunetti et al., 2001; Gitti et al., 2002), or secondary electrons produced by hadronic collisions (e.g. Dennison, 1980; Blasi & Colafrancesco, 1999). Relics, on the other hand, are arc-shaped structures extending up to 2 Mpc and are found more often at the outskirts of clusters. They are believed to originate from collisionless shocks that accelerate cosmic-ray electrons (Blandford & Eichler, 1987; Ryu et al., 2003), or from the reacceleration of pre-existing fossil cosmic-ray electrons (e.g. Markevitch et al., 2005; van Weeren et al., 2016a), which are relic particles from past energetic events.

More recently, the first diffuse radio bridges connecting two pre-merging galaxy clusters have been discovered as well (Govoni et al., 2019; Botteon et al., 2020b). This is a direct detection at radio frequencies of filamentary structure between galaxy clusters, as part of the large-scale structure of the cosmic web. Images with the well-known pre-merging clusters Abell 399 and Abell 401 with corresponding radio bridge in the middle are shown in Figure 1.2. The origin of the non-thermal radio emission from radio bridges between pre-merging clusters has been the subject of debate in recent years, but the latest studies seem to favour a turbulent reacceleration model (Brunetti & Vazza, 2020; Nunhokee et al., 2021; Pignataro et al., 2024), similar to the most favoured mechanism for generating radio haloes. Some of the currently open questions related to radio bridges are: Is turbulent reacceleration the only and best explanation for radio bridges? How common are (large-scale) radio bridges? What is their connection with the evolution of the large-scale structure

of our Universe? To answer these questions, very deep observations are required at low frequencies, preferably combined with X-ray observations to compare thermal and non-thermal emissions, since the reacceleration of electrons occurs in similar regions as where the ICM is heated (e.g. Brunetti & Jones, 2014). We explore this in Chapter 2.

1.3. Catching radio waves

Although radio astronomy began with Jansky’s work at very low frequencies, using equipment originally developed for early radio communication, the field later shifted its focus to higher frequencies, as more advanced telescopes were designed and built at these frequencies. Among these are for instance the Westerbork Synthesis Radio Telescope (WSRT Baars et al., 1973), the Very Large Array (VLA; Thompson et al., 1980), or the more recently constructed Atacama Cosmology Telescope (ALMA; Swetz et al., 2011). At higher frequencies, astronomers have the opportunity to study radio emissions from molecular clouds (e.g. Yusef-Zadeh, 2012), protoplanetary disks (e.g. Villenave et al., 2020), star-forming regions (e.g. Rosero et al., 2016), among other astronomical objects which emit thermal emission.

We already hinted that detecting non-thermal radio emission, such as associated with the jets discussed in Section 1.2.1 or the radio bridges mentioned in Section 1.2.2, requires telescopes specifically designed to observe low-frequency radio waves. Some low-frequency telescopes, like for example the Cambridge Interferometer (Ryle & Hewish, 1955) or the Ukrainian T-shaped Radio Telescope (UTR; Braude et al., 1978),² were operational for many decades. However, these instruments were limited by their resolution. This is because the resolution improves as frequency increases for a given telescope size, according to the following relation

$$\theta \propto \frac{\lambda}{D_{\text{tel}}}, \quad (1.1)$$

where θ is the resolution, λ is the wavelength, and D_{tel} the diameter of the telescope (or baseline as we will see in the next subsection). This implies that improving the ability of radio telescopes to resolve more objects and substructures at low frequencies requires more advanced designs and techniques.

²Sadly, this telescope was, during my PhD, heavily damaged during the Russian invasion of Ukraine in 2022.




Figure 1.2: Image of pre-merging clusters Abell 399-401 at optical (Pan-STARRS; Chambers et al., 2016), X-ray (XMM-Newton; Sakelliou & Ponman, 2004), and detected at radio frequencies at 144 MHz (de Jong et al., 2022). Each galaxy cluster – one in the north and the other in the south – exhibits a radio halo, which largely overlaps with X-ray emission. This indicates that these clusters have experienced a galaxy cluster merger in the past. At the same time, there is also fainter emission between the galaxy clusters represented by a radio bridge, showing that these clusters are also in the process of merging with each other. Radio galaxies with their optical counterparts are also depicted. This image utilises the data analysed in Chapter 2 and was designed by F. Sweijen.

1.3.1. Interferometry

A method to improve resolution at lower frequencies is through interferometry,³ which is a technique where signals received by two or more widely spaced dishes or antennas are correlated to extract information from objects on smaller angular scales (e.g. Ryle & Vonberg, 1946; Kellermann & Moran, 2001; Quirrenbach, 2009). This relies on the principle of interference, where radio waves from the same source in the sky, arriving at different antennas, are combined in such a way that their constructive or destructive interference patterns at different angular scales are used to extract information about the source. The correlated signals between antenna pairs, called ‘visibilities’, are represented by complex numbers (with a phase and amplitude) and correspond to baselines, which represent the distances between each pair of antennas or dishes. These visibilities with their corresponding baseline coordinates (commonly expressed in the uv -plane or uvw -plane), are samples of the Fourier transforms of the sky brightness. This indicates that they can be mapped into image space using (fast) Fourier transforms (FFT; Cooley & Tukey, 1965). In this way, interferometry effectively creates a virtual telescope with a maximum diameter equal to the longest baseline, enabling significantly higher resolution than any single-dish telescope could achieve on its own. In this context, D_{tel} in Equation 1.1 refers to the length of the longest baseline.

The challenge for interferometers, compared to single-dish telescopes, is their limited coverage of the uv -plane, which affects the image reconstruction process. A straightforward way to improve this coverage is through making use of Earth’s rotation, a technique known as aperture synthesis (Ryle & Hewish, 1960). This allows telescopes to observe the same region of the sky from different angles over time, effectively sampling different parts of the uv -plane at various frequencies. The incomplete sampling of the uv -plane introduces also artefacts, typically from side-lobes in the point spread function (PSF), which can distort or obscure (fainter) features in the image. Mitigating this effect is typically achieved during the conversion of visibilities into more accurate images using (variations of) the CLEAN algorithm (Högbom, 1974; Clark, 1980). This algorithm iteratively deconvolves the radio image by identifying and subtracting point sources starting with the initial ‘dirty’ image, progressively reconstructing the true sky brightness distribution by adding the subtracted components in each iteration to a model image. The final components in the model image are then convolved with a Gaussian beam and added back to the residual noise background, resulting in a cleaner image with reduced side-lobe interference. This algorithm can be further extended with various enhancements, such as multi-scale CLEAN, which accounts for emission structures at multiple spatial scales (e.g. Cornwell, 2008).

³Note that this technique can be used at non-radio wavelengths as well.

1.3.2. The Low Frequency Array

Since the early 2000s, advancements in technology, financial resources, and renewed interest have enabled the development of state-of-the-art low-frequency interferometers with longer baselines and higher sensitivities. These include the Giant Metrewave Radio Telescope (GMRT; Swarup et al., 1991), the Murchison Widefield Array (MWA; Tingay et al., 2013), the Long Wavelength Array (LWA; Ellingson et al., 2009), the upcoming Square Kilometer Array (SKA; Carilli & Rawlings, 2004), and the International Low Frequency Array (LOFAR) Telescope (ILT; van Haarlem et al., 2013). LOFAR is the primary instrument that serves us in the next two chapters of this thesis to address some of the outstanding scientific questions discussed in Sections 1.2.1 and 1.2.2, motivating us later in the thesis to explore ways to enhance its capabilities.

LOFAR is a unique instrument that detects low frequency radio waves between 10 and 80 MHz with its low-band antennas (LBAs) and between 110 and 240 MHz with its high-band antennas (HBAs). It makes use of phased-array technology (e.g. Herd & Conway, 2016), where signals from many individual dipole antennas are combined electronically. This allows LOFAR to ‘steer’ its focus across the sky by introducing delays in the data, without physically moving its antennas. The absence of moving parts, combined with LOFAR’s focus on detecting long wavelengths, reduces the need for strict mechanical engineering tolerances compared to telescopes designed for higher frequencies or with moving parts. This enables a cost-effective design and construction of a large number of antennas throughout Europe. LOFAR has a dense concentration of core and remote stations in the Netherlands and international stations across many European countries. The current longest baseline stretches 1980 km between Birr (Ireland) and Łazy (Poland). An overview of the station layout across Europe is presented in Figure 1.3.

The size of the field of view (FoV) of a LOFAR station is given by

$$\text{FoV} \propto \left(\frac{\lambda}{D_{\text{stat}}} \right)^2,$$

where D_{stat} is the station diameter. This relation implies that LOFAR is able to obtain a large FoV in the order of multiple square degrees. In fact, the FoV achievable using only the Dutch stations at 150 MHz is more than 10 deg^2 (e.g. Shimwell et al., 2017), while an international station is larger and therefore reaches at the same frequencies, combined with the Dutch stations, a FoV of about 6 deg^2 (e.g. Sweijen et al., 2022c). LOFAR’s capability to image large FoVs makes it a great instrument to perform surveys over large sky areas, catching many sources at once.




Figure 1.3: LOFAR stations across Europe. The orange circles indicate stations currently active and from which data are utilised in this thesis, with the Dutch LOFAR core stations represented by a larger circle. The blue circles mark LOFAR stations currently being constructed and soon to be added to the international LOFAR network. The purple circle is the Kilpisjärvi Atmospheric Imaging Receiver Array (KAIRA; McKay-Bukowski et al., 2015), which is a LOFAR station in Finnish Lapland that operates independently of the main LOFAR network. We have marked the currently active longest baseline of 1980 km between Birr and Łazy with a red dashed line.

1.3.3. Sensitivity vs. resolution with LOFAR

The Dutch LOFAR stations provide a $6''$ resolution and a sensitivity of approximately $70 \mu\text{Jy beam}^{-1}$ for one 8 hours observation (Shimwell et al., 2019). By incorporating the international stations, the resolution can improve to $0.3''$, along with an improved sensitivity by about a factor of 2, due to the larger collecting area of the international stations (van Haarlem et al., 2013; Varenus et al., 2015). This sensitivity improvement refers to the point source sensitivity, as imaging at lower resolutions still performs better in tracing extended low surface brightness emission from large-scale objects (e.g. radio haloes, relics, or bridges), which may otherwise be hidden in the background noise in high-resolution images. Intermediate arcsecond-scale resolution offers therefore a valuable alternative as well, striking a balance between resolution, sensitivity, and computational cost (e.g. Ye et al.,




Figure 1.4: Demonstrating the power of resolving sources at higher resolution. Both panels are from left to right postage stamp images from the following maps: 6'' resolution radio map from (Shimwell et al., 2025), 0.3'' resolution radio map from de Jong et al. (2024), and optical map from DESI (Dey et al., 2019) with 0.6'' radio contours from de Jong et al. (2024) on top. *Upper panel:* Although the 6'' resolution is sensitive to capture low surface brightness emission, it does in this case not allow for accurate morphological classifications of RLAGN in FRI or FRII types, while on 0.3'' resolution the galaxies are highly resolved. *Lower panel:* While at 6'' the sources are difficult to classify and there is no clear optical counterpart, do the high-resolution counterparts show that these are individual bent RLAGN jets.

2024).

For the studies of RLAGN, low surface brightness sensitivity is advantageous for detecting diffuse jets of FRIs when the object is resolved and classifiable at a given resolution (e.g. Mingo et al., 2019). Yet, in cases such as presented in Figure 1.4, where we compare detections of neighbouring RLAGN at 6'' and 0.3'', it is clear that while some of the RLAGN are detected at 6'' resolution, they are not enough resolved to successfully classify these sources. Jurlin et al. (2024) recently illustrated the benefit of using the multi-resolution ability of LOFAR, by revealing recent activity of RLAGNs close to their core with the international long baselines, while uncovering older emissions from past activity using the Dutch shorter baselines. In the context

of galaxy clusters, van Weeren et al. (2024) demonstrated that combining multiple resolutions from LOFAR offers a more comprehensive understanding of the various processes occurring in and around galaxy clusters as well. While these examples illustrate the need to balance resolution and sensitivity for specific science cases, it is also crucial to emphasise that processing the data to produce these images does not come for free, as will be outlined in the next section.

1.4. Data processing with LOFAR

The combination of high resolution, high sensitivity, and large FoV positions LOFAR as a powerful instrument for producing deep, high-resolution wide-field images of the low-frequency radio sky. However, collecting so much information in a single observation also comes at a cost. Hence, efficient and robust processing pipelines that calibrate and image the data are important to efficiently map the raw visibility data to a final science-ready image.

1.4.1. Calibration at low frequencies

For standard data processing, several effects need first to be addressed. These include calibration for different systematic effects, such as clock offsets, polarisation offsets, ionospheric effects, beam errors, and Faraday rotation (e.g. van Weeren et al., 2016b; de Gasperin et al., 2019a; Mevius, 2018). At low frequencies, the main calibration challenges are coming from direction-dependent effects (DDEs). These are variations in observed signals caused by changes in direction across the FoV of a telescope. At the low frequencies observed by LOFAR, these variations are particularly introduced by the ionosphere. This is a layer in Earth’s atmosphere filled with charged particles that interact with incoming radio waves, especially at low frequencies. These interactions cause phase delays, which vary across the sky and change over time, leading to image distortions and positional shifts. Multiple software packages have been developed to deal with this issue, such as DP3⁴ (van Diepen et al., 2018a; Dijkema et al., 2023), SPAM⁵ (Intema et al., 2009), Sagecal⁶ (Kazemi et al., 2011), KillMS⁷ (Tasse, 2014a,b; Smirnov & Tasse, 2015), and facetselfcal⁸ (van Weeren et al., 2021). To account for the DDE variation across the FoV, the corrections from these packages are typically applied within a facet-based approach. This requires the selection and calibration of a number of calibrator sources across

⁴<https://dp3.readthedocs.io>

⁵<http://www.intema.nl/doku.php?id=huibintema:spam:start>

⁶<https://github.com/nlesc-dirac/sagecal>

⁷<https://github.com/saopicc/killMS>

⁸https://github.com/rvweeren/lofar_facet_selfcal

the FoV with enough signal at the largest baselines. The field is then divided into facets using a Voronoi tessellation (Schwab, 1984; van Weeren et al., 2016b; Tasse et al., 2018), wherein the calibration corrections are applied to smaller sky areas (the facets) corresponding to the nearest calibrator source.

1.4.2. Imaging large FoV

Having done all the calibration, the corrected visibilities need to be converted to images using Fourier transforms (typically in the form of FFTs). Although LOFAR’s large FoV is an advantage to capturing many sources at once, it also implies that it is not possible to estimate the sky to be a flat plane, as for imaging smaller FoVs. So, to incorporate curvature of the Earth and the sky, it is essential to account for the so-called ‘ w -term’, which describes the non-coplanarity of an interferometer (e.g. Perley, 1999; Cornwell, 2008). This additional term makes applying Fourier transforms computationally more expensive. Fortunately, several efficient imaging software packages are around to effectively take into account the curvature to generate wide-field images. One of these is **WSClean**⁹ (Offringa et al., 2014), which uses w -stacking to group visibilities with similar w -values into different layers in the Fourier domain. Another software package is **DDFacet**¹⁰ (Tasse et al., 2018), which uses w -kernels to correct for the sky curvature within facets.

1.5. Pipelines and computational cost

For LOFAR data with only Dutch stations, the DDF-pipeline¹¹ was developed (Shimwell et al., 2019; Tasse et al., 2021). This pipeline calibrates and processes visibilities automatically up to a fully DDE-corrected wide-field image at 6'' resolution. This has resulted in the LOFAR Two-metre Sky Survey (LoTSS; Shimwell et al., 2017, 2019, 2022; Williams et al., 2019), which has nearly completed imaging the entire northern hemisphere. Yet, a fully automated pipeline for direction-dependent calibration and wide-field imaging data that includes both Dutch and international LOFAR stations does currently not exist, despite having a general outline of the necessary steps (Morabito et al., 2022a, 2025a; Sweijen et al., 2022c; Ye et al., 2024; de Jong et al., 2024). This is because reducing this data has been experimental until now, and the high computational cost associated with full data processing must first be managed before scaling up.

⁹<https://wsclean.readthedocs.io/en/latest>

¹⁰<https://github.com/saopicc/DDFacet>

¹¹<https://github.com/mhardcastle/ddf-pipeline>

1.5.1. Computational Bottlenecks

Prior to the work presented in this thesis, the computational cost for reducing a single 8 hours LOFAR observation for sub-arcsecond wide-field imaging summed up to about $\sim 250,000$ CPU core hours (Sweijen et al., 2022c).¹² This large cost is due to the substantial visibility data volumes, which for one LOFAR observation is about 16 terabytes (TB). Using data compression techniques this can be brought down to about 4 TB (Offringa, 2016), or alternatively being averaged with baseline-dependent averaging (BDA; Cotton, 1986, 2009; Wijnholds et al., 2018).¹³ This data volume cannot be further reduced, since it is essential to contain the full time and frequency resolution for the final imaging to avoid time and bandwidth smearing. In particular, the final imaging, despite having efficient software (as outlined in Section 1.4.2), has proven to be the main computational bottleneck, since this is where the large number of visibilities with calibration solutions applied have to be iteratively converted to about 8 billion image pixels (Sweijen et al., 2022c; de Jong et al., 2024). Given these large computational costs associated with the processing of LOFAR data for sub-arcsecond wide-field imaging, it is essential to optimise our algorithms and data reduction strategies for greater efficiency.

1.5.2. CORTEX

As part of the Centre for Optimal Real-Time Machine Studies of the Explosive Universe (CORTEX) project,¹⁴ we aim to address the computing and automation challenges for processing LOFAR data for deep sub-arcsecond wide-field imaging. First of all, our goal is to further automate the processing of LOFAR high-resolution data, reducing the need for human intervention. This includes automatically selecting the best calibrator sources and parameters for correcting DDEs during LOFAR data reduction but also validating the selection and calibration. This can for instance be done by finding ways to assess the signal-to-noise (S/N) for different calibrators at the longest baselines directly from visibility data but also by incorporating state-of-the-art neural networks to assess image quality post-calibration and fine-tune parameters accordingly, as will be discussed in Chapter 4. Second of all, to reduce the computational costs of processing large datasets, it is important to revisit the current data reduction strategy and develop techniques to reduce the data volume before imaging, as will be discussed in Chapters 5 and 6.

¹²This means that it would take about 7 years to process this data on a typical 4-core laptop.

¹³Dysco compression does currently not work in combination with BDA. However, alternative compression techniques are currently being explored and may be able to combine lossy or loss-less data compression with BDA compression (e.g. Dodson et al., 2024).

¹⁴<https://www.nwo.nl/en/projects/nwa116018316>

1.5.3. Wide and deep surveys

By developing an automatic robust pipeline and by bringing down the computational cost, not only surveys across large sky areas (similar to LoTSS) are enabled but also ultra-deep imaging at high resolutions, deeper than what can be achieved at $6''$ resolution. At $6''$ resolution confusion noise – additional background noise caused by blending faint, unresolved sources (e.g. Condon et al., 2012) – becomes a limiting factor for reaching sensitivities below $10 \mu\text{Jy beam}^{-1}$ (e.g. Sabater et al., 2021). Having beams 400 times smaller in area with LOFAR’s $0.3''$ resolution will therefore mitigate this issue. For instance, having hundreds of hours of integration time already available of the ELAIS-N1 deep field in the LOFAR long-term archive (LTA), gives us the opportunity to produce an ultra-deep 2.5×2.5 degrees image with sensitivities in the order of $\mu\text{Jy beam}^{-1}$.

1.6. This thesis

In this thesis, we aim to leverage LOFAR’s exceptional sensitivity (depth) and resolution (detail) to study our Universe at low frequencies, while at the same time pushing its capabilities to open new windows for scientific exploration. We begin by demonstrating LOFAR’s capability for creating deep sensitive wide-field images of large-scale structures through the production of the deepest image of a known radio bridge, aiming to explain its origin. We then utilise at smaller scales the resolution and sensitivity of the Dutch LOFAR stations with the goal of answering open questions about the cosmic evolution of extended RLAGN morphologies. Although this leads to new insights, it highlights the resolution limitations of Dutch-only observations as well, motivating us to enhance data processing techniques for creating deep sub-arcsecond wide-field images in the second half of this thesis. This effort results in the deepest sub-arcsecond wide-field image to date, along with new techniques to push our work further with reduced need for human intervention while saving computational cost. The chapters in this thesis are outlined as follows:

- In **Chapter 2**, we apply an advanced calibration technique to create the deepest wide-field image of a known radio bridge using LOFAR data. This enables us to investigate its particle (re)acceleration mechanisms and explore its origin. Through a point-to-point comparison of radio and X-ray data, our findings indicate that turbulent (second-order Fermi) reacceleration is the most likely cause of this phenomenon. We suggest a similar origin for the radio haloes of the individual clusters as well.
- In **Chapter 3**, we study the cosmic evolution of FRI and FRII radio galaxies up to redshift $z = 2.5$ by selecting classified FRI and FRII radio sources

detected by LOFAR. This provides us with 100 times more FRIs above $z = 0.3$ compared to previous work. With radio luminosity functions we find space density enhancements towards higher redshifts for the most powerful radio galaxies, while we observe a space density decrease over redshift for less powerful sources. The former can be physically explained by a higher gas density in an earlier universe, while the latter is likely related to selection biases that can only be overcome by improving the spatial resolution and depth of our observations.

- In **Chapter 4**, we create the deepest sub-arcsecond 2.5×2.5 degrees image with LOFAR of the ELAIS-N1 deep field with a central RMS noise of about $14 \mu\text{Jy beam}^{-1}$. This is accomplished by combining four different observations of the same pointing centre, enhancing the direction-independent calibration, and further automating the direction-dependent calibrator selection. While computing costs are reduced, they remain significant, underscoring the need to minimise these costs for deeper imaging or wide surveys.
- In **Chapter 5**, we address a large fraction of the cost associated with deep imaging by revisiting and improving the technique we call ‘sidereal visibility averaging’. This method makes use of the fact that baseline coordinates repeat every sidereal day, which allows us to average visibilities at similar baseline coordinates when combining multiple observations. This reduces for instance the computing costs for imaging by a factor of 10 when producing an ultra-deep image with 500 hours of LOFAR data. Although the method has proven to be reliable, we also stress the fact that it is important to take into account precession and aberration, which introduce offsets in the baseline coordinates between observations.
- In **Chapter 6**, we address remaining bottlenecks that prevent the processing of a large number of observations of the same sky area for sub-arcsecond wide-field imaging with LOFAR. This includes image quality improvements, the introduction of a streamlined framework, a refined and further automated DDE calibration using neural networks, and the lowering of the total computing costs up to a factor of 4 compared to the work from Chapter 4. This work lays the foundation for processing and imaging hundreds of hours of existing LOFAR data of the ELAIS-N1 field, ultimately enabling the creation of a wide-field sub-arcsecond image with unprecedented depth in the order of a few $\mu\text{Jy beam}^{-1}$.

1.7. Future prospects

LOFAR will remain for at least another decade a leading instrument to study radio frequency emission below 200 MHz. This is because its capability to achieve sub-arcsecond resolution will remain unmatched at these frequencies, even with the advent of the SKA (Carilli & Rawlings, 2004). As already mentioned in Section 1.5.3, having both high-resolution and high sensitivity will be the main way to beat confusion noise and probing the depths of our Universe. The work presented in this thesis, contributes to this promising future for LOFAR, as we continue to process and analyse increasing amounts of LOFAR data.

1.7.1. Towards a better understanding of AGN and galaxy clusters

Scientifically, LOFAR’s current and upcoming deep (sub-)arcsecond data allows us to better separate radio emission due to AGN and star formation, even within the same galaxy, using brightness temperature measurements (e.g. Morabito et al., 2022b, 2025b). Furthermore, high-resolution views of more nearby radio galaxies give us the ability to unravel the detailed physics and dynamics of their jets as well (e.g. Mahatma et al., 2023; Mahatma, 2023; Jurlin et al., 2024). Moreover, sub-arcsecond spatial detail enables the possibility to more accurately measure radio galaxy sizes over large sky areas (e.g. Sweijen et al., accepted), helping us for example to enhance our understanding of the evolution of their jets on a more statistical basis. Similarly, we can with enough deep sub-arcsecond resolution data also correct for the selection biases identified in Chapter 3, and obtain a deeper understanding of how radio jet morphologies relate to the cosmic evolution at redshifts beyond $z = 0.8$. At larger scales, deep high-resolution images of the sky at 144 MHz may also provide unprecedented insights into emissions from various scales in galaxy clusters, from thin filaments constituting the tails of radio galaxies to diffuse threads embedded in the ICM (De Rubeis et al. in prep.).

1.7.2. LOFAR 2.0

The upcoming upgrade to LOFAR 2.0 will further improve LOFAR’s capabilities. One of the main highlights of this upgrade is the synchronisation of the clocks of the Dutch stations. This will mitigate clock drifts for these stations and therefore improve calibration (de Gasperin et al., 2019b). The LOFAR 2.0 upgrade will also enhance the sensitivity in the LBA and enable simultaneous LBA and HBA observations (e.g. Edler et al., 2021). This will benefit studying objects at wider frequency bandwidth at the same time (spectral studies), but may also help to improve LBA high-resolution imaging (e.g. Morabito et al., 2016; Groeneveld et al., 2022), by uti-

lizing combined calibration at HBA and LBA to solve for the ionospheric-induced delays.

1.7.3. FAIR community-wide collaboration

As the complexity of our data reduction strategies increases, it will also be essential to unite efforts across the LOFAR community and across multiple disciplines. Tackling new challenges efficiently will require a collaborative approach, where skilled software developers play a key role in advancing the data reduction pipelines and making those more robust, while astronomers contribute to the development and help steer focus on achieving the scientific goals. These synergies will require thorough communication as data reduction strategies need to be optimised and computational bottlenecks need to be addressed, but the primary goal to push scientific boundaries should never be lost out of sight. Adhering also to FAIR principles – being Findable, Accessible, Interoperable, and Reusable – will make it easier for other researchers to find, access, and build upon future data and tools (Wilkinson et al., 2016; O’Toole & Tocknell, 2022), accelerating advancements in LOFAR sub-arcsecond imaging and studying the Universe across multiple wavelengths and domains.

1.7.4. Technological advancements

Looking further ahead, we can expect more implementations of code and algorithms for the use of GPUs, pushing the efficiency of our data processing. And artificial intelligence (AI) will not only play a bigger role in our daily lives but also continue to enhance our ability to process data and deepen our understanding of the Universe. Likewise, emerging technologies such as quantum computing may hold the potential to revolutionise how we live and work as well. However, despite the promises of these advancements, it is crucial to approach their implementation critically and carefully. Ethical considerations (e.g. Kop et al., 2023; Resnik & Hosseini, 2024) and energy consumption (e.g. Kruithof et al., 2023) must remain important as we improve our software and integrate new technologies into scientific research and beyond. Ultimately, the success of innovative solutions depends not only on their capabilities but also on the ability to use them wisely.