



Universiteit
Leiden
The Netherlands

Facets of radio-loud AGN evolution : a LOFAR surveys perspective

Williams, W.L.

Citation

Williams, W. L. (2015, December 10). *Facets of radio-loud AGN evolution : a LOFAR surveys perspective*. PhD Thesis. Retrieved from <https://hdl.handle.net/1887/37016>

Version: Not Applicable (or Unknown)

License: [Leiden University Non-exclusive license](#)

Downloaded from: <https://hdl.handle.net/1887/37016>

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

Cover Page



Universiteit Leiden



The handle <http://hdl.handle.net/1887/37016> holds various files of this Leiden University dissertation.

Author: Williams, Wendy L.

Title: Facets of radio-loud AGN evolution : a LOFAR surveys perspective

Issue Date: 2015-12-10

4 LOFAR 150 MHz observations of the Boötes field: Catalogue and Source Counts

We present the first wide area (19 deg^2), deep ($\approx 120 - 150 \mu\text{Jy beam}^{-1}$), high resolution ($5.6 \times 7.4 \text{ arcsec}$) LOFAR High Band Antenna image of the Boötes field made at $130 - 169 \text{ MHz}$. This image is at least an order of magnitude deeper and $3 - 5$ times higher in angular resolution than previously achieved for this field at low frequencies. The observations and data reduction, which includes full direction-dependent calibration, are described here. We present a radio source catalogue containing 5 652 sources detected over an area of 19 deg^2 , with a peak flux density threshold of 5σ . As the first thorough test of the facet calibration strategy, introduced by van Weeren, we investigate the flux and positional accuracy of the catalogue. We present differential source counts that reach an order of magnitude deeper in flux density than previously achieved at these low frequencies, and show flattening at 147 MHz flux densities below 10 mJy associated with the rise of the low flux density star-forming galaxies and radio-quiet AGN.

Williams, W. L., van Weeren, R. J., Röttgering, H. J. A., et al.
MNRAS, submitted

4.1 Introduction

The LOw Frequency ARray (LOFAR) is a new generation radio telescope operating at 10–240 MHz (van Haarlem et al. 2013). Its large instantaneous field-of-view, combined with multi-beaming capabilities, high-spatial resolution, and large fractional bandwidth make LOFAR an ideal instrument for carrying out large surveys of the sky which will have long-lasting legacy value. As such, ‘Surveys’ is one of the six LOFAR Key Science Projects (KSP). The science goals of the Surveys KSP are broad, covering aspects from the formation and evolution of large-scale structure of the Universe; the physics of the origin, evolution and end-stages of radio sources; the magnetic field and interstellar medium in nearby galaxies and galaxy clusters; to Galactic sources. The deep LOFAR surveys will be crucial in the study of AGN evolution and the history of black-hole accretion. In particular, the Surveys KSP aims to answer questions related to the nature of the different accretion processes, the properties of the host galaxies, the role of AGN feedback in galaxy growth and evolution, the radio-source duty cycle and the relation of the AGN with their environment (e.g. Heckman & Best 2014, and references therein). The radio-source population has not been well-studied at low flux densities and low frequencies. To achieve the diverse goals of the LOFAR surveys, which will be carried out over the next five years, a tiered approach is being used: Tier-1 covers the largest area at the lowest sensitivity, and will include low-band (LBA; 15 – 65 MHz) and high-band (HBA; 110 – 180 MHz) observations across the whole 2π steradians of the northern sky with a targeted rms noise of ≈ 0.1 mJy beam $^{-1}$ and a resolution of ≈ 5 arcsec. Deeper Tier-2 and Tier-3 observations will cover smaller areas, focussing on fields with the highest quality multi-wavelength datasets available (for details see Rottgering 2010).

Several low-frequency surveys have been performed in the past, such as the Cambridge surveys 3C, 4C, 6C and 7C at 159, 178, 151 and 151 MHz, respectively (Edge et al. 1959; Bennett 1962; Pilkington & Scott 1965; Gower et al. 1967; Hales et al. 1988, 2007), the UTR-2 sky survey between 10 – 25 MHz (Braude et al. 2002), and the VLSS at 74 MHz (Cohen et al. 2007; Lane et al. 2014). The GMRT significantly improved low frequency imaging, particularly in terms of sensitivity and angular resolution, and several GMRT surveys have now been performed at 153 MHz (e.g. Ishwara-Chandra & Marathe 2007; Sirothia et al. 2009; Ishwara-Chandra et al. 2010; Intema et al. 2011), 325 MHz (e.g. Mauch et al. 2013) and 610 MHz (e.g. Garn et al. 2007, 2008a,b). Recently, the Murchison Widefield Array (MWA; Lonsdale et al. 2009; Tingay et al. 2013), operating at 72 – 231 MHz, has yielded the GaLactic and Extragalactic All-sky MWA survey (GLEAM; Wayth et al. 2015). GLEAM covers the entire Southern sky ($\delta < 25^\circ$) with a noise level of a few mJy beam $^{-1}$ and angular resolution of a few arcminutes. However, for extragalactic science in particular, LOFAR provides a significant advantage in both image resolution and sensitivity.

Advanced calibration and processing techniques are needed to obtain deep high-fidelity images at low radio frequencies. In particular, direction-dependent effects (DDEs) caused by the ionosphere and imperfect knowledge of the station beam shapes need to be corrected for. van Weeren (prep, submitted) have recently presented a new scheme for calibrating the direction-dependent effects and imaging LOFAR data that combines elements from existing direction-dependent calibration methods such as SPAM (Intema et al. 2009) and SAGECAL (Yatawatta et al. 2013; Kazemi et al. 2011). The Boötes field observations presented here also serve as a testbed for this calibration strategy, which allows us to produce science quality images at the required Tier-1 survey depth.

Here we report on the first LOFAR Cycle 2 High Band Antenna (HBA) observations of the Boötes field. The Boötes field is one of the Tier-3 Survey fields and the aim is to eventually

survey this field to the extreme rms depth of $12 \mu\text{Jy beam}^{-1}$ (1σ) at 147 MHz. The Boötes field has been extensively studied at higher radio frequencies and in other parts of the electromagnetic spectrum. Radio observations have been carried out at 153 MHz with the GMRT, both as a single deep 10 deg^2 pointing (Intema et al. 2011) and as a seven-pointing 30 deg^2 mosaic (Williams et al. 2013). Further observations include those at 325 MHz with the VLA (Croft et al. 2008; Coppejans et al. 2015), and deep, $28 \mu\text{Jy beam}^{-1}$ rms, 1.4 GHz observations with WSRT (de Vries et al. 2002). The field has also been observed with the LOFAR Low Band Antennae at 62 MHz (van Weeren et al. 2014).

The Boötes field is one of the largest of the well-characterised extragalactic deep fields and was originally targeted as part of the NOAO Deep Wide Field Survey (NDWFS; Jannuzi et al. 1999) covering $\approx 9 \text{ deg}^2$ in the optical (B_W , R , I) and near infra-red (K) bands. There is a wealth of ancillary data available for this field, including X-ray (Murray et al. 2005; Kenter et al. 2005), UV (GALEX; Martin et al. 2003), and mid infrared (Eisenhardt et al. 2004). The AGN and Galaxy Evolution Survey (AGES) has provided redshifts for 23, 745 galaxies and AGN across 7.7 deg^2 of the Boötes field (Kochanek et al. 2012). This rich multiwavelength dataset, combined with the new low frequency radio data presented here, will be important for determining the evolution of black-hole accretion over cosmic time.

The outline of this paper is as follows. In Section 4.2 we describe the LOFAR observations covering the NOAO Boötes field. In Section 4.3 we describe the data reduction techniques employed to achieve the deepest possible images. Our data reduction relies on the ‘Facet’ calibration scheme (van Weeren prep, submitted) which corrects for direction-dependent ionospheric phase corruption as well as LOFAR beam amplitude corruption. In Section 4.4 we present the final image and describe the source-detection method and the compilation of the source catalogue. This section also includes an analysis of the quality of the catalogue. The spectral index distribution and differential source counts are presented in Section 4.5. Finally, Section 4.6 summarises and concludes this work. Throughout this paper, the spectral index, α , is defined as $S_\nu \propto \nu^\alpha$, where S is the source flux density and ν is the observing frequency. We assume a spectral index of -0.8 unless otherwise stated.

4.2 Observations

The Boötes field was observed on 2014 August 10 with the LOFAR High Band Antenna (HBA) stations. An overview of the observations is given in Table 4.1. By default, all four correlation products were recorded with the frequency band divided into 195.3125 kHz-wide subbands (SBs). Each SB was further divided into 64 channels. The integration time used was 1 s in order to facilitate the removal of radio frequency interference (RFI) at high time resolution. The maximum number of SBs for the system in 8 bit mode is 488 and the chosen strategy was to use 366 for the Boötes field giving a total bandwidth of 72 MHz between 112–181 MHz. The remaining 122 SBs were used to observe the nearby bright calibrator source, 3C 294, located 5.2° away, with a simultaneous station beam, with SBs randomly spread between 112–181 MHz. The main observations were preceded and succeeded by 10 min observations of the primary flux calibrators 3C 196 and 3C 295, respectively, with identical SB setup to the Boötes observation, i.e. 366 SBs (72 MHz bandwidth) between 112–181 MHz. For the observations 14 Dutch remote and 24 Dutch core stations were used. This setup results in baselines that range between 40 m and 120 km. The uv -coverage for the Boötes field observation is displayed in Fig. 4.1. The ‘HBA_DUAL_INNER’ configuration was employed. In this configuration, the core stations are

Table 4.1: LOFAR HBA observation parameters.

Observation IDs	L240762 (3C 196) L240764 (Boötes, 3C 294) L240766 (3C 295)
Pointing centres	08 : 13 : 36 48 : 13 : 03 (3C 196) 14 : 32 : 00 +34 : 30 : 00 (Boötes) 14 : 06 : 44 +34 : 11 : 25 (3C 294) 14 : 11 : 20 52 : 12 : 10 (3C 295)
Integration time	1 s
Observation date	2014 August 10
Total on-source time	10 min (3C 196, 3C 295) 8 hr (Boötes, 3C 294)
Correlations	XX, XY, YX, YY
Sampling mode	8-bit
Sampling clock frequency	200 MHz
Frequency range	112–181 MHz
Bandwidth	71.48 MHz (Boötes, 3C 196, 3C 295) 23.83 MHz (3C 294)
Subbands (SBs)	366 contiguous (Boötes, 3C 196, 3C 295) 122 randomly distributed ^a (3C 294)
Bandwidth per SB	195.3125 kHz
Channels per SB	64
Stations	62 total 14 remote 24 core (48 split)

^abetween 112–181 MHz

each split into two substations (48 total), and only the inner ≈ 30.8 m of the remote stations (which have a total diameter of 41 m) are used to obtain similar station beam sizes to the core stations (which have a diameter of 30.8 m). The resulting half-power beam width (HPBW) is $\approx 4.2^\circ$ ¹ at 150 MHz.

4.3 Data Reduction

In this section we describe the calibration method that was used to obtain the required deep high-fidelity high-resolution images. The data reduction and calibration consists of two stages: a non-directional and a directional part to correct DDEs caused by the ionosphere and imperfect station beam models. The non-directional part includes the following steps:

1. initial flagging and removal of RFI;

¹Based on the calculated average primary beam for the Boötes observation (see Section 4.3.2)




Figure 4.1: uv -coverage for the Boötes field at 130 – 169 MHz. The maximum baseline is 120 km (or $60 k\lambda$). Only one out of every ten uv -points in time and one out of every 40 points in frequency are plotted: the plot nevertheless shows how the large fractional bandwidth fills the uv -plane radially. The two colours show the symmetric uv points obtained from the conjugate visibilities.

2. solving for the calibrator complex gains, including ‘clock-TEC (Total Electron Content) separation’, and transfer of the amplitudes, median clock offsets and the XX – YY phase offset from calibrator to the target field;
3. removal of bright off-axis sources;
4. averaging; and
5. amplitude and phase (self-)calibration of the target field at medium (20–30 arcsec) resolution

This is then followed by a scheme to correct for DDEs in order to reach near thermal-noise-limited images using the full resolution offered by the longest ‘Dutch-LOFAR’ baselines of about 120 km. All calibration steps are performed with the `BLACKBOARD SELFCAL` (BBS) software system (Pandey et al. 2009) and other data handling steps were undertaken with the LOFAR Default Pre-Processing Pipeline (DPPP). These steps are explained in more detail below. The direction-dependent calibration scheme is described in full by van Weeren (prep, submitted).

We used the full frequency coverage of the 10 min 3C 196 observation as the primary calibrator observation to derive the time-independent instrumental calibration including amplitudes, median clock offsets and the XX – YY phase offset. We did not use the simultaneous, but sparse, frequency coverage on the calibrator 3C 294. For the Boötes field we selected 200 out of the total 366 observed SBs (55 per cent) covering the frequency range 130–169 MHz for further processing. The main limitation to the number of subbands processed was computational time; the main data reduction was carried out on a node with 20 virtual cores on ASTRON’s CEP3

cluster¹ and most of the parallelisation in the reduction is performed over 10-SB blocks, so that the full ‘facet’ reduction could be achieved in a reasonable time.

4.3.1 Direction-independent calibration

Flagging and RFI removal

The initial preprocessing of the data was carried out using the Radio Observatory pipeline and consisted of RFI excision (using AOFlogger; Offringa et al. 2010, 2012), flagging the noisy first channel channel and last three channels of each SB, and averaging in time and frequency to 2 s and 8 channels per SB. The data were stored at this resolution in the LOFAR Long Term Archive (LTA) at this point. One core station (CS007) and one half of another core station (CS501HBA1) were flagged entirely due to malfunction (failure to record data or low gains). CS013 was also flagged entirely due its different design (the dipoles are rotated by 45°).

Calibration transfer from Primary Calibrator 3C 196

Using BBS we obtained parallel hand (XX and YY) gain solutions for 3C 196, on timescales of 2 s for each frequency channel independently. In this step we also solved simultaneously for ‘Rotation Angle’ per station per channel to remove the effects of differential Faraday Rotation from the parallel hand amplitudes. The solutions were computed with the LOFAR station beam applied to separate the beam effects from the gain solutions. We used a second-order spectral model for 3C 196 consisting of 4 point sources separated by 3 – 6 arcsec each with a spectral index and curvature term (V. N. Pandey (ASTRON), priv. comm.). This is close to, but not exactly the same as the Scaife & Heald (2012) flux scale. We used these calibration solutions to determine the direction-independent and time-invariant instrumental calibrations, including amplitude calibration, correction of clock delays between the remote and core stations and an offset between the XX and YY phases.

The remote LOFAR stations have their own clocks which are not perfectly synchronized with the single clock that is used for all the core stations – the offsets between these clocks can be of the order of 100 ns, which is large enough to cause strong phase delays within a single SB for the remote-remote and core-remote baselines. In addition to these clock offsets, which appear to be similar from observation to observation, the clocks can drift by ≈ 15 ns over a few hours. The primary calibrator phase solutions were used to calculate the clock offsets to be applied to the target observation (we do not correct for the 15 ns clock drifts). We used the ‘Clock-TEC separation’ method described in detail by van Weeren (prep, submitted) that uses the frequency-dependent phase information across the full frequency range to separate the direction-independent clock errors from the direction-dependent ionospheric effects. The clock phase errors, or delays, vary linearly with frequency (phase $\propto \delta t \times \nu$, where δt is the clock difference), while the ionospheric phases vary inversely with frequency (phase $\propto \text{dTEC} \times \nu^{-1}$, where dTEC is the differential Total Electron Content). Fitting was performed on a solution interval timescale of 5 s and smoothed with a running median filter with a local window size of 15 s. The results are shown in Fig. 4.2, from which it can be seen that the clock offsets for the remote stations in our observation were between -90 ns and 80 ns with respect to the core. The ionospheric conditions during the 10 min calibrator observation were good, showing relatively smooth variations in the differential TEC of $\approx 0.2 \times 10^{16} \text{ m}^{-2}$.

¹Each node has 2 ten-core Intel Xeon e5 2660v2 (25M Cache, 2.20 GHz) processors with 128 GB RAM




Figure 4.2: Fitted clock and TEC differences from the 3C196 phase solutions. The fitting was performed on the solution interval timescale of 2 s and smoothed with a running median filter with a local window size of 10 s. Top: Fitted station clock offsets as a function of time with respect to the reference (core) station CS001HBA0. The core stations all have clock values close to zero, while the remote stations have large clock offsets with respect to the core stations, between -90 ns and 80 ns. Bottom: Differential TEC values for the same stations. These show typical dTEC values (≈ 0.2 TEC units). The core stations all have similar dTEC values around zero as they essentially probe the same patch of ionosphere. RS310 has the largest variation in dTEC and also has discrete jumps in its dClock values as a result of its noisier amplitudes.

For a few stations we found small but constant offsets between the XX and YY phases. We determined these offsets by taking the median phase difference between the XX and YY phases during the 3C 196 observation for each station.

The amplitudes were inspected for outliers and smoothed in the frequency axis with a running median filter with window size of 3 SBs (≈ 0.6 MHz), and a single median value in the time axis of the 10 min observation.

These calculated median clock offsets, XX–YY phase offsets and amplitude values were transferred from the 10 min observation of 3C 196 to the target-field data. The resulting target-field visibilities have amplitudes in Janskys and are free of global clock offsets (note that we do not correct for the 15 ns clock drifts during the target observation) and XX–YY phase offsets.

Removal of bright off-axis sources

A few radio sources are sufficiently bright to contribute flux through the sidelobes of the station beams, the amplitudes of which are strongly modulated in frequency, time and baseline as they move in and out of the station-beam sidelobes. To remove these effects we simply predicted the visibilities of the brightest of these ‘A-team’ sources (Cyg A, Cas A, Vir A, and Tau A) with the station beam applied in BBS, and flagged all times, frequencies and baselines where the contributed flux density from these sources exceeded 5 Jy. The amount of data flagged in this step was typically 2–5 per cent per SB.

Averaging

The data were then averaged in time to a more manageable 8 s and two channels per SB (98 kHz channelwidth). The main limit on the time resolution is set by the requirement to avoid decorrelation due to rapid ionospheric phase variations. Even at this time and frequency resolution we expect some smearing at large radial distances from the field centre due to time and bandwidth smearing, of the order of 7 and 10 per cent respectively at the half-power point of the primary beam (2.1° from the pointing centre) and 13 and 15 per cent respectively at 2.5° at 150 MHz. The individual corrected SBs were combined in groups of ten, providing datasets of ≈ 2 MHz bandwidth and 20 channels, each of which was ≈ 30 GB in size.

Self-calibration of target

We used a single 10-SB block at 148–150 MHz on which to perform direction-independent self-calibration. We started with a model derived from the 30 deg^2 GMRT image of the Boötes field at 153 MHz (Williams et al. 2013). Then a Gaussian component model was derived using PyBDSM (Mohan & Rafferty 2015), which identifies islands of contiguous emission above a given island threshold (PyBDSM parameter `thresh_is1`) around pixels exceeding a given pixel threshold (`thresh_pix`). The thresholds are given in factors above the varying background rms determined within a sliding box of a given size (`rms_box`), which takes into account the increased rms near bright sources due to calibration artefacts. The source detection parameters are described in more detail by Williams et al. (2013). The resolution of this initial sky model is 25×25 arcsec – the native resolution of the GMRT image – and it includes all sources in the GMRT image greater than 20 mJy ($\approx 6\sigma$). The brightest 10 sources in the GMRT model were replaced by Gaussian components taken from the ≈ 5 arcsec resolution FIRST catalogue (Becker et al. 1995), with their total GMRT 150 MHz flux density and flux density ratios taken from FIRST.




Figure 4.3: Final self-calibrated image for the 10-SB block at 148–150 MHz. The resolution is 23×20 arcsec. The greyscale shows the flux density from -3σ to 25σ where $\sigma = 1 \text{ mJy beam}^{-1}$ is the approximate rms noise in the central part of the image. Calibration artefacts are clearly visible around the brightest sources as only direction-independent self-calibration has been performed.

The self-calibration was performed with two iterations with phase-only solutions followed by two iterations with amplitude and phase solutions. The solution interval in the first phase-only iteration was 8 s and we obtained a single solution in frequency, neglecting phase changes within the 2 MHz bandwidth. The signal-to-noise ratio within the 10-SB block was high enough to obtain good-quality, i.e. coherent, solutions in each timestep. The station beam model was applied in the ‘solve’ step in BBS and the solutions were applied. The data were corrected for the station beam in the phase centre. Imaging was carried out using AWIMAGER (Tasse et al. 2013), which deals with the non-coplanar nature of the array and performs a proper beam correction across the field-of-view in imaging. From the resulting ‘corrected’ images we created a new sky model, again using PyBDSM (with `rms_box = (boxsize, stepsize) = (150, 40)` pixels, `thresh_pix = 5 σ` and `thresh_is1 = 3 σ`), and included all Gaussians in the model. We then performed an additional iteration of phase-only self-calibration followed by two iterations of phase and amplitude self-calibration before making a final self-calibrated image for the 10-SB block at 149 MHz. A part of this image is shown in Fig. 4.3. The noise level achieved is ≈ 1 mJybeam. While deeper images at a similar resolution can be made with more bandwidth, this is sufficient for the purpose of initial calibration prior to the direction-dependent calibration. All the imaging at this stage was carried out with a field-of-view of 6.4° at a resolution of ≈ 20 arcsec by imposing an outer uv -limit of 40 k λ and Briggs (1995) robust weighting (`robust = 0`). This weighting results in a slightly lower resolution, with less emphasis on the calibration artefacts. The imaging at each step performed in two stages: first without a mask and then with a mask. The CLEAN masks for each image were generated automatically from the first image (without a mask) by using PyBDSM (with `rms_box = (85, 30)` pixels) to create a 3σ island threshold map.

All the 10-SB blocks were then corrected for the station beam in the phase centre before further processing. We then used the final image of the self-calibration cycle at 148 – 150 MHz to make a Gaussian input sky model with PyBDSM to (self)-calibrate the remaining 10-SB blocks. All the components in this model were assumed to have a spectral index of -0.8 . For all the bands we performed a single phase and amplitude calibration against this model and made images of each band with the Common Astronomy Software Applications (CASA; McMullin et al. 2007) version 4.2.1 using W -projection (Cornwell et al. 2008, 2005) to handle the non-coplanar effects; CASA does not allow for the full beam correction across the field-of-view but it does allow a much larger field to be imaged than with AWIMAGER. Because the beam correction was not performed in imaging, the resulting images are *apparent* sky images. The imaging was carried out in two iterations. The first image was made at ‘medium resolution’ using a uv -cut of 7 k λ and Briggs weighting (`robust = 0`) to limit the resolution to ≈ 20 arcsec, with 7.5 arcsec pixels; the field-of-view imaged was $\approx 10^\circ$. Next, in each 10-SB dataset we subtracted the CLEAN components with their self-calibration gain solutions and re-imaged the subtracted data using a ‘low resolution’ of ≈ 2 arcmin with a uv -cut of 2 k λ and Briggs weighting (`robust = 0`), 25 arcsec pixels, and a field-of-view of $\approx 30^\circ$. This field-of-view extends to the second sidelobe of the station beams, allowing us to image and subtract these sources. The low-resolution image also picks up extended low-surface-brightness emission in the field not CLEANed in the medium-resolution image. Both medium- and low-resolution images were created with CLEAN masks automatically generated from images made without CLEAN masks and PyBDSM with `rms_box = (50, 12)` and `rms_box = (60, 12)` respectively. The CLEAN components in the low-resolution image were then subtracted in the same way as the medium-resolution components and a combined list of CLEAN components for the medium- and low-resolution images was created.

The resulting products are 20 sets of 2-MHz residual datasets between 130 and 169 MHz – i.e. with all sources out to the second sidelobes subtracted using the gain solutions, but with the

residual data itself not corrected for the gain solutions. In addition, there are 20 corresponding CLEAN-component *apparent* sky models of the sources that were subtracted. These products serve as the input for the direction-dependent calibration scheme, which is described in the following section.

4.3.2 Directional calibration – ‘Facet’ scheme

Significant artefacts remain in the self-calibrated images, even at 20 arcsec resolution, and the rms noise of a few mJy beam^{-1} is a factor of 3–5 higher than what is expected with these imaging parameters. Both of these issues result from the direction-dependent effects of the station beams and ionosphere. To correct for these we follow the direction-dependent calibration, or ‘facet’, scheme, of van Weeren (prep, submitted), in which calibration is performed iteratively in discrete directions and imaging is carried out within mutually exclusive facets around each calibration direction. The key concern is keeping the number of degrees of freedom in the calibration small with respect to the number of measured visibilities, in order to facilitate solving for DDEs in tens of directions. The ‘facet’ scheme has the following underlying assumptions:

1. the only calibration errors are a result of ionosphere and beam errors;
2. the station beams vary slowly with time and frequency;
3. differential Faraday rotation is negligible, so that XX and YY phases are affected identically by the ionosphere;
4. the phase frequency dependence is $\text{phase} \propto \nu^{-1}$ as a result of ionosphere only; and
5. DDEs vary slowly across the field-of-view.

For a more detailed description and discussion of the underlying problems and assumptions see van Weeren (prep, submitted) and references therein. The following sub-sections describe our implementation of the ‘facet’ scheme on the Boötes field data. At this point we selected 20 contiguous 10-SB datasets covering the frequency range of 130–169 MHz, i.e. 55 per cent of the total observed bandwidth. This selection was made primarily as a trade-off between the expected achievable depth and computational constraints.

Facetting the sky

Initially we identified 28 calibration directions or groups consisting of single bright sources or closely (few arcminute) separated sources with a combined total flux density $\gtrsim 0.3 \text{ Jy}$. The bright source positions or centres of the groups define the facet directions that were used to tile the sky using Voronoi tessellation, ensuring that each point on the sky lies within the facet of the nearest calibrator source. The assumption here is that the calibration solutions for the calibrator group applies to the full facet. The typical facet size is a few tens of arcminutes in diameter. Around and beyond the HPBW of the station beam, the assumption that DDEs, in particular the beam, vary slowly across the field-of-view breaks down. We found that two of the facets initially defined were too large and showed worsening calibration artefacts away from the calibrator source. These two facets were subdivided into smaller facets each having new calibrator groups. The final set of facets is shown in Fig. 4.4. The following steps were then performed for each direction sequentially, starting with the brightest calibrator sources.




Figure 4.4: Facet coverage of the Boötes field. The grey boxes show the positions of the calibrator directions and the size of these boxes show the area used for self-calibration, i.e. larger boxes include more sources. The coloured polygons show the Voronoi tessellation of the image plane based on these calibrator positions. The maximum size of the facets is limited to 50 arcmin radius from the calibrator direction (2048 pixels at $1.5 \text{ arcsec pixel}^{-1}$) resulting in some incomplete coverage particularly outside the FWHM. The black circle has a radius of 2.44° , at the approximate 40 per cent power point of the average primary beam.

Directional self-calibration

The CLEAN components of all the sources within the calibrator group were added back in each 10-SB dataset (with the direction-independent calibration solutions with which they had been initially subtracted). These datasets were then phase-rotated to the direction of the calibration group and averaged in frequency (but not in time) to 1 channel per 2-MHz dataset. The frequency averaging greatly speeds up calibration without any bandwidth-smearing effects in the small (few arcminute) calibrator group images.

A self-calibration cycle with four iterations was then performed. The imaging at each iteration was carried out using CASA with the full 130 – 169 MHz bandwidth and multi-frequency synthesis (MFS) CLEAN with a second-order frequency term (i.e. including spectral index and curvature, $n_{\text{terms}} = 2$) (Conway et al. 1990), with the automated masking described in Section 4.3.1. This resulted in a CLEAN-component model with both flux and spectral index. Multi-scale CLEAN (MS-MFS; Rau & Cornwell 2011; Cornwell 2008) was used only in the case of a few complex extended sources. We used a Briggs robust parameter of -0.25 , a pixel size of 1.5 arcsec, and imposed a uv -minimum of 80λ to achieve a resolution of approximately 5.6×7.4 arcsec. Using a more negative robust weighting allows for higher resolution images after the direction-dependant effects have been accounted for.

In the first two self-calibration cycles we solved for a single Stokes I phase-offset and TEC term per station in groups of 5 10-SB datasets, i.e. within 10 MHz bands. This gave 8 parameters (4 phase and 4 TEC) per station per solution interval. The solutions were computed on timescales of 8 s where possible, but this was increased to 16 s for fainter calibration groups and even to 24 s for the faintest calibration groups where the signal-to-noise was lower.

After having solved for, and applied, the short timescale phase-offset and TEC solutions, we then ran a third and fourth iteration of self-calibration in which we solved for XX and YY gains independently per 10-SB dataset on timescales between 5 and 20 min, depending on the flux density of the calibrator group. This yielded an additional 4 parameters per 10-SB block. This calibration primarily takes out the slowly-varying complex station beams. As expected the phase component of these solutions is small because the fast phase component has been taken out. In the final self-calibration cycle, the amplitudes were normalised to unity across the full frequency range to prevent changes to the flux-scale. The normalisation corrections were typically smaller than a few percent. Example solutions are shown in Fig. 4.5 for a single direction and in Fig. 4.6 for all directions as a snapshot in time. Both the amplitudes and phases show consistent values between directions. When viewed as a movie in time, trends can be seen ‘moving across’ the field-of-view, in particular in the fast phases, which is consistent with ionospheric phase disturbances propagating through the field-of-view. The significant improvement made in the self-calibration cycles is demonstrated by the calibrator images shown in Fig. 4.7. It is clear that both the phase *and* amplitude calibration are currently required: the phase distortions resulting from ionospheric effects will always have to be corrected for, but future improvements in the LOFAR beam models may eliminate the need for the amplitude calibration.

Facet imaging and subtraction

After obtaining the direction-dependent calibration solutions for the given direction, the remaining sources within the facet were added back (with the direction-independent solutions with which they had been subtracted). Assuming the solutions for the calibrator group apply to the full facet, we applied those solutions to the facet data at the original 2 channels per SB




Figure 4.5: Example good DDE solutions for a few selected stations obtained for a single direction s3 (the images corresponding to this direction are shown in Fig. 4.7). Left: The effective Stokes I phase corrections, evaluated at an arbitrary frequency of 150 MHz. The solutions are obtained on a timescale of 8 s using 10 MHz of bandwidth. Centre and Right: the additional XX and YY amplitude (centre) and phase (right) solutions for the 160 – 162 MHz SB block obtained on a timescale of 10 min after application of the short-timescale phase offsets and TEC solutions. In all cases phases are plotted with respect to core station CS001HBA0.




Figure 4.6: Example good DDE solutions for a few selected stations obtained for all directions for a given timestep. Top: XX amplitude solutions for the 160 – 162 MHz SB block obtained within a particular 10 min time interval. Bottom: The effective Stokes I phase corrections evaluated at an arbitrary frequency of 150 MHz and plotted as an angle. In both plots, the two outer and inner circles show, respectively, the 30 and 50 per cent power points of the average station beam. The solutions are obtained on a timescale of 8 s using 10 MHz of bandwidth. Phases are again plotted with respect to core station CS001HBA0.




Figure 4.7: Images showing the improvements during the DDE calibration for a few example directions. All images are made using the full dataset (120–181 MHz, $n_{\text{terms}} = 2$, $\text{robust} = -0.25$) and have a resolution of 5.6×7.4 arcsec. Note that at this resolution many of the bright DDE calibrator sources are resolved. The leftmost column shows the initial images made with only the direction-independent self-calibration solutions. The centre column displays the improvements after two iterations of TEC and phase-offset-only DDE calibration step. The right column shows the improvement after the final DDE self-calibration iteration after two further iterations of DDE TEC+phase offset calibration and DDE XX and YY gain (amplitude and phase). For all four directions shown, the TEC+phase offsets were solved for on 8 s time intervals. The XX and YY gains were solved for on 10 min timescales. The scalebar in each image is 1 arcmin.

resolution, which allows the $1/\nu$ dependence of the TEC term to be applied on a channel-to-channel basis. At this point the corrected facet data were averaged 5 times in frequency and 3 times in time (to 0.5 MHz per channel and 24 s) to avoid bandwidth and time smearing within the facet image. Each facet was then imaged with MS-MFS CLEAN with `nterms` = 2 over the full 130 – 169 MHz bandwidth, with a Briggs robust parameter of -0.25 , 1.5 arcsec pixels and a uv -minimum of 80λ . Note that since all facets have different phase centres, their uv coverage differs slightly and so the restoring beams are slightly different. As in Section 4.3.1 we do not use AWIMAGER mainly due to its limitations in imaging beyond the HPBW of the station beam.

Each facet image provided an updated sky model that was then subtracted from the full-resolution data with the corresponding direction-dependent solutions, thereby improving the residual data to which the subsequent facets were added. This process (from Section 4.3.2) was repeated until all facets had been calibrated and imaged. The order in which the facets were handled is determined by the severity of the calibration artefacts in the direction-independent images, which roughly corresponds to the brightness of the calibration groups, so that the directions with the worst artefacts were corrected first and did not influence later directions. After all the facets were calibrated in this way, we re-imaged all the facets. This step removes the artefacts present in the given facet resulting from bright sources in neighbouring facets which had only been calibrated after the given facet. It reduces the rms by a few per cent. This was achieved by adding back the facet sky model to the residual data, applying the relevant directional-dependent solutions, and imaging with the same parameters (using the full 130 – 169 MHz bandwidth with `nterms` = 2, `robust` = -0.25 , a pixel-size of 1.5 arcsec, and a uv -minimum of 80λ). At this point we applied a common restoring beam of 5.6×7.4 arcsec to all the facets; this resolution was chosen as the smallest beam containing all the fitted restoring beam from the individual facets.

The resulting products are (i) high-resolution images of the facets, combined in a single image (‘mosaic’) covering the full field-of-view, (ii) short-timescale corrections for the phases variation of the ionosphere, and (iii) long-timescale phase and amplitude corrections for the station beams in the directions of each facet.

‘Primary beam’ correction

CASA was used to image the individual facets so the images are in ‘apparent’ flux units, with the station beam taken out only in the phase centre. In order to measure real sky fluxes we performed a primary-beam correction to take into account the LOFAR station beam response across the field-of-view. We used AWIMAGER to produce an average primary beam map from all the 20 10-SB datasets, using the same Briggs weighting (`robust` = -0.25), but with much larger pixels (5.5 arcsec), and imaging out to where the average station beam power drops to 20 per cent (approximately 6.2° in diameter). We corrected the combined ‘mosaic’ image by dividing by the regridded AWIMAGER average primary beam map¹. We imposed a ‘primary beam’ cut where the average primary beam power drops below 40 per cent, at an approximate radius of 2.44° , which results in an image covering a total area of $\approx 19 \text{ deg}^2$.

4.4 Final Image and Catalogue

The combined ‘mosaic’ image at 5.6×7.4 arcsec resolution is shown in Fig. 4.8. The central rms noise level is relatively smooth and $\lesssim 125 \mu\text{Jy beam}^{-1}$, and 50 per cent of the map is at

¹The actual average primary beam map used is the square root of the AWIMAGER output.

a noise level below $180 \mu\text{Jy beam}^{-1}$ (see also Fig. 4.11). A small portion of the image covering the inner 0.25 deg^2 is shown in Fig. 4.9 to illustrate the resolution and quality of the map. There remain some phase artefacts around the brightest sources (see for example the source in the lower right of the image in Fig. 4.9), which have not been entirely removed during the facet calibration; however, these are localised around the bright sources. Fig. 4.10 shows a comparison between the 153 MHz GMRT image, the LOFAR 148 – 150 MHz direction-independent self-calibration image (see Section 4.3.1) and the final 130 – 169 MHz direction-dependent calibrated image for three arbitrary positions. This serves to illustrate the significant improvement in both noise and resolution achieved with the new LOFAR observations over the existing GMRT data, which has an rms noise level of $\approx 3 \text{ mJy beam}^{-1}$ and resolution of 25 arcsec.

4.4.1 Source Detection and Characterisation

We compiled a source catalogue using PyBDSM (Mohan & Rafferty 2015) to detect and characterise sources. We ran PyBDSM separately on each facet image, using the pre-beam-corrected image as the detection image and the primary beam-corrected image as the extraction image. The rms map was determined with a sliding box `rms_box = (160, 50)` pixels (i.e. a box size of 160 pixels every 50 pixels), with a smaller box `rms_box = (60, 15)` pixels in the regions around bright sources (defined as having peaks exceeding 150σ). Using a smaller box near bright sources accounts for the increase in local rms as a result of calibration artefacts. For source extraction we used `thresh_pix = 5\sigma` and `thresh_isl = 3\sigma` (i.e. the limit at which flux is included in the source for fitting). Fig. 4.11 illustrates the variation in rms noise thus determined across the combined ‘mosaic’ image.

PyBDSM fits each island with one or more Gaussians, which are subsequently grouped into sources. We used the `group_tol` parameter with a value of 10.0 to allow larger sources to be formed. This parameter controls how Gaussians within the same island are grouped into sources; the value we chose is a compromise between selecting all Gaussians in a single island as a single source, thus merging too many distinct nearby sources, and selecting them as separate sources, thus separating the radio lobes belonging to the same radio source. Sources are classified as ‘S’ for single sources and ‘M’ for multiple-Gaussian sources. PyBDSM reports the fitted Gaussian parameters as well as the deconvolved sizes, computed assuming the image restoring beam. Uncertainties on the fitted parameters are computed following Condon (1997). The total number of sources detected by PyBDSM in all the facets is 5 652 comprised of 10 771 Gaussian components of which 3 010 were single-component sources. We allowed PyBDSM to include sources that were poorly fitted by Gaussians; these sources are included in the catalogue with the integrated flux density being the total flux density within the source island and flagged as having poor Gaussian fits. Additionally, based on visual inspection, 49 sources were flagged as artefacts near bright sources, or false detections, or detections on the edge of the image. We also identified and flagged sources lying on the facet boundaries as these lie in two facets and may have corrupted fluxes because their CLEAN components were added back in separate facet images.

4.4.2 Astrometric and Flux Density Uncertainties

In this section we evaluate the uncertainties in the measured LOFAR flux densities and positions. For both the flux densities and positions we make corrections to the catalogue we present in Section 4.4.5, to account for systematic effects.

Given the uncertainties in the low-frequency flux density scale (e.g. Scaife & Heald 2012)




Figure 4.8: Greyscale map showing the entire mosaic. The image covers $\approx 19 \text{ deg}^2$. The greyscale shows the flux density from $-3\sigma_{\text{cen}}$ to $25\sigma_{\text{cen}}$ where $\sigma_{\text{cen}} = 110 \mu\text{Jy beam}^{-1}$ is the approximate rms in the mosaic centre.




Figure 4.9: Zoom-in of the central part of the mosaic. The image covers 0.25 deg^2 . The greyscale shows the flux density from $-3\sigma_{\text{cen}}$ to $25\sigma_{\text{cen}}$ where $\sigma_{\text{cen}} = 110 \mu\text{Jy beam}^{-1}$ is the approximate rms in the mosaic centre.




Figure 4.10: Images showing a comparison between a few random sources in the GMRT image (*left* column), the LOFAR 148 – 150 MHz direction-independent self-calibration image (*centre* column, see Section 4.3.1) and the final 130 – 169 MHz direction-dependent calibrated image (*right* column). The noise in the three images is respectively ≈ 3 , ≈ 1 and $\approx 0.15 \text{ mJy beam}^{-1}$ and the resolution is 25 , 20 and $5.6 \times 7.4 \text{ arcsec}$ respectively. Each image is 15 arcmin on a side. The image centre J2000 coordinates (right ascension and declination) are shown in degrees in the top left of each row and they are all plotted on the same colourscale.




Figure 4.11: Left: Greyscale map showing the local rms noise measured in the mosaic image. The greyscale shows the rms noise from $0.5\sigma_{\text{cen}}$ to $4\sigma_{\text{cen}}$, where $\sigma_{\text{cen}} = 110 \mu\text{Jy beam}^{-1}$ is the approximate rms in the mosaic centre. The contours are plotted at $125 \mu\text{Jy beam}^{-1}$ and $175 \mu\text{Jy beam}^{-1}$. Peaks in the local noise coincide with the locations of bright sources. Right: Cumulative area of the map with a measured rms noise level below the given value.

and the LOFAR station beam models, we may expect some systematic errors in the measured LOFAR flux densities. In order to determine any systematic offsets we have compared the LOFAR flux densities to those of the GMRT image of the Boötes field at 153 MHz (Williams et al. 2013), selecting only sources detected at high signal-to-noise ($S_{\text{peak}}/\sigma > 10$) in both maps. For this sub-sample of 448 sources we determined the ratio of integrated flux densities between LOFAR and the GMRT $f_S = S_{\text{LOFAR}}/S_{\text{GMRT}}$. We measured a median ratio of $\langle f_S \rangle = 1.04$ with a standard deviation of $\sigma_{f_S} = 0.22$. The flux density ratio showed no significant variation with distance from the phase centre. However we noticed a small variation with position on the sky, plotted in Fig. 4.12. The variation is such that the LOFAR flux densities are slightly higher towards the North-North-West (upper right of the map) and lower towards the South-South-East (lower left of the map). We note that the trend is consistent across different facets so is a result of some global effect and not a result of discrete calibration in facets. Moreover, the trend is similar, but noisier, if we consider LOFAR sources extracted from the self-calibrated-only image from the single 10-SB block at 148–150 MHz. This further suggests that the flux offset is not introduced by the direction dependent calibration procedure. The flux density ratio is plotted as a function of the position angle between the source and the phase centre in Fig. 4.13, to which we fitted a sinusoid of the form $f_S = 1.05 + 0.10 \sin(116^\circ + \phi)$. The median ratio was $\langle f_S \rangle = 1.00$ with a standard deviation of $\sigma_{f_S} = 0.20$ after applying this sinusoid function to correct the LOFAR flux densities. This correction assumes that the GMRT flux densities are ‘correct’. To evaluate this we performed a similar comparison with sources in the deep 1.4 GHz WSRT catalogue of the Boötes Field (de Vries et al. 2002). Scaling the higher frequency flux densities to 150 MHz with a spectral index of -0.8 we recovered a similar, albeit noisier, trend for the sample of high signal-to-noise LOFAR sources matched to the WSRT sources. This suggests that it is likely not to be GMRT pointing errors causing systematic trends in the GMRT flux densities. We speculate that the observed trend in the flux density errors may be the result of an incorrect primary-beam correction, which itself may be caused by errors in the LOFAR station beam model used in AWIMAGER, or a map projection error. Finally, we assume an overall flux density error of 15 per cent, to account for systematic effects (see also Intema et al. 2011; Williams et al. 2013).

It is possible that errors in the phase calibration can introduce uncertainties in the source positions. We used the 1.4 GHz FIRST catalogue (Becker et al. 1995) to assess the LOFAR positional accuracy. The uncertainty on the FIRST positions is given by $\sigma = s(1/\text{SNR} + 0.05)$, where s is the source size in arcsec. FIRST counterparts were identified within 2 arcsec of our LOFAR sources with high signal-to-noise ratios, i.e. $S_{\text{peak}}/\sigma_{\text{local}} > 10$. This yielded 968 matches, of which 313 were fitted with a single Gaussian in the LOFAR image, making them more likely to be point sources. For the high signal-to-noise point source sample, we measured a small mean offset in right ascension of $\Delta\alpha = \alpha_{\text{LOFAR}} - \alpha_{\text{FIRST}} = -0.062 \pm 0.001$ arcsec, with rms $\sigma_\alpha = 0.43$ arcsec and a somewhat larger mean offset in declination of $\Delta\delta = \delta_{\text{LOFAR}} - \delta_{\text{FIRST}} = -0.321 \pm 0.002$ arcsec with rms $\sigma_\alpha = 0.55$ arcsec. The offset is negligible and we note that it is within the size of the LOFAR image pixels (1.5 arcsec). However, closer inspection of the offsets showed that the offset in declination varied systematically across the full 5° field-of-view, and to a lesser extent for the offsets in right ascension. A correction for this offset has been made by fitting a plane, $\Delta = a\alpha + b\delta + c$, where Δ is in arcsec, to the offset values and applying the fitted offsets to all sources in the catalogue. While this could be expressed as a rotation and therefore a sinusoidal correction could be made, as we have done for the flux densities, we find that the positional offsets are asymmetric with respect to the pointing centre. The fitted planes were $\Delta\alpha = -0.12(\alpha - 218^\circ) + 0.03(\delta - 34.5^\circ) - 0.03$ and $\Delta\delta = 0.13(\alpha - 218^\circ) + 0.30(\delta + 34.5^\circ) - 0.36$. These offsets




Figure 4.12: Map of the measured ratios of integrated flux densities for high signal-to-noise LOFAR sources with respect to the GMRT. The colourscale shows the the flux density ratio.




Figure 4.13: Measured ratios of integrated flux densities for high signal-to-noise LOFAR sources with respect to the GMRT as a function of source position angle with respect to the phase centre (plotted as points). The colourscale shows the distance of each source to the pointing centre. The blue curve shows a fitted sinusoid which we used for correcting the positions of the LOFAR catalogued sources. The solid and dashed horizontal lines show the measured median and standard deviation of the flux density ratio.

and functional corrections are shown in Fig. 4.14. After applying this correction we measured $\Delta\alpha = \alpha_{\text{LOFAR}} - \alpha_{\text{FIRST}} = -0.017 \pm 0.001$ arcsec with rms deviation $\sigma_{\Delta\alpha} = 0.39$ arcsec, and $\Delta\delta = \delta_{\text{LOFAR}} - \delta_{\text{FIRST}} = -0.007 \pm 0.001$ arcsec, with $\sigma_{\Delta\delta} = 0.38$ arcsec. We note that since the initial phase calibration of the LOFAR data was performed against the GMRT model sky, these positional offsets may originate from the GMRT model. We did evaluate the GMRT positions in a similar way and found no significant variations, but this may be due to the much larger positional uncertainty and the 25 arcsec synthesised beam of the GMRT map. We also consider LOFAR sources extracted from the self-calibrated-only image from the single 10-SB block at 148–150 MHz compared to FIRST and find a similar trend. This suggests that the position offsets are not introduced by the direction dependent calibration procedure. We hypothesise that a map projection or mosaic construction error could be the cause of these offsets, but note that it is a small correction.

The scatter in the offsets between the LOFAR and FIRST positions is a combination of noise-independent calibration errors, ϵ , in both the LOFAR and FIRST data as well as a noise-dependent error, σ , from position determination via Gaussian-fitting:

$$\sigma^2 = \epsilon_{\text{LOFAR}}^2 + \epsilon_{\text{FIRST}}^2 + \sigma_{\text{LOFAR}}^2 + \sigma_{\text{FIRST}}^2$$

To separate the noise-dependent and -independent uncertainties we select from the above sample only the FIRST sources with position errors of less than 0.6 arcsec and measure an rms scatter of $(\sigma_\alpha, \sigma_\delta)_{\text{LOFAR}} = (0.37 \text{ arcsec}, 0.35 \text{ arcsec})$ between the corrected LOFAR and FIRST source positions for this very high signal-to-noise sub-sample of 89 sources. From Becker et al. (1995), the FIRST calibration errors are $(\epsilon_\alpha, \epsilon_\delta)_{\text{FIRST}} = (0.05 \text{ arcsec}, 0.05 \text{ arcsec})$. The noise-dependent fit errors for both the LOFAR and FIRST can be assumed to be small so we determine the LOFAR calibration errors to be $(\epsilon_\alpha, \epsilon_\delta)_{\text{LOFAR}} = (0.37 \text{ arcsec}, 0.35 \text{ arcsec})$. This scatter may contain a small contribution resulting from any spectral variation between 150 and 1400 MHz on scales smaller than the resolution of the surveys (≈ 5 arcsec).

4.4.3 Reliability

The reliability of a source catalog indicates the probability that all sources above a given flux density are real. It is defined as the fraction of all detected sources in the survey area above a certain total flux density limit that are real sources and are not accidental detections of background features or noise. To estimate the reliability, we extracted sources from the inverted residual mosaic image, assuming that negative image background features are statistically the same as positive ones. We grouped the detected negative ‘sources’ by total flux density into 20 logarithmic flux density bins and compare these to the binned results of the regular (positive) source extraction as described in Section 4.4.1. For convenience, we define the real number of sources to be the number of positive sources minus the number of negative sources.

The left panel of Fig. 4.15 shows the false detection rate, FDR , determined from the number ratio of negative sources over positive sources per flux density bin. The right panel shows the reliability curve, $R = 1 - FDR$, determined from the number ratio of real sources over positive sources above the total flux density limits that define the lower edges of the flux density bins. For a 1 mJy total flux density threshold, the reliability is 90 per cent.

4.4.4 Resolved Sources

In the absence of noise, resolved sources can easily be identified based on the ratio of the integrated flux density to the peak flux density, $S_{\text{int}}/S_{\text{peak}} > 1$. However, since the uncertainties

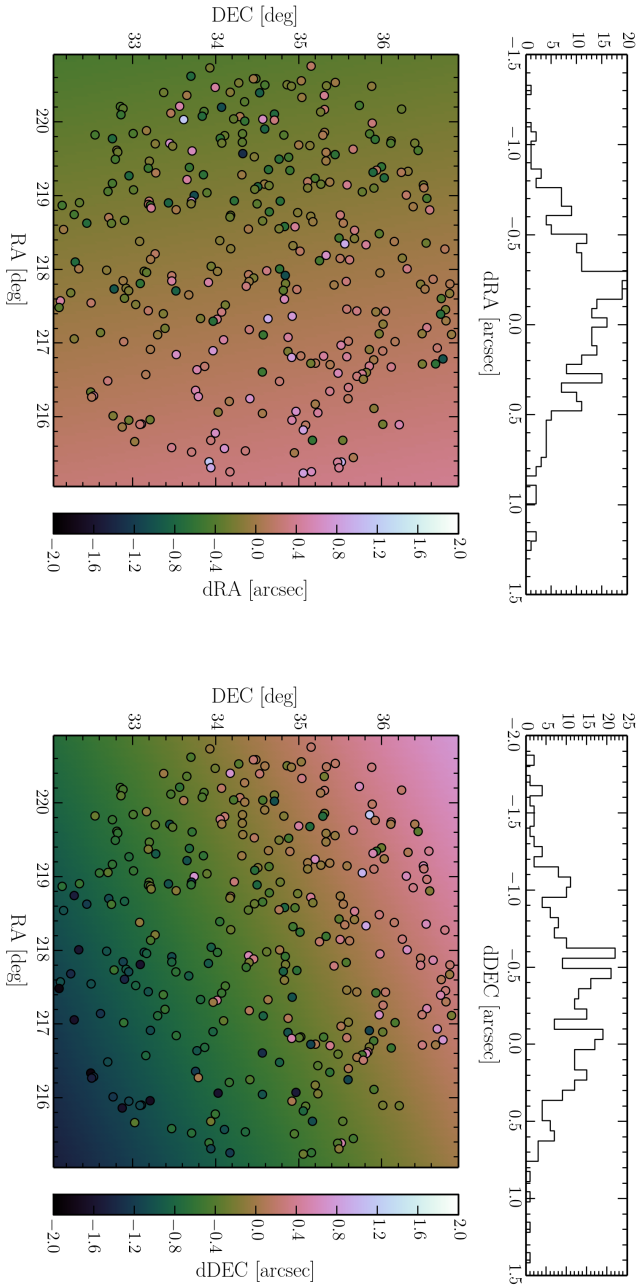


Figure 4.14: Measured offsets in right ascension (left) and declination (right) for high signal-to-noise sources of LOFAR with respect to FIRST (plotted as points). The colourscale shows the value of the given offset. The top panels show the distribution of the offsets. The offsets, particularly in declination, show a variation across the field-of-view. The background colourscale shows a fitted plane used for correcting the LOFAR catalogued sources.




Figure 4.15: *Left* False detection rate as a function of peak flux density to local signal-to-noise ratio calculated from source detection on an inverted residual map. *Right* Estimated reliability of the catalogue as a function of integrated flux density limit accounting for the varying sensitivity across the field-of-view.

on S_{int} and S_{peak} are correlated, the $S_{\text{int}}/S_{\text{peak}}$ distribution is skewed, particularly at low signal-to-noise. The flux density ratio as a function of signal-to-noise is shown in Fig. 4.16. We note that the scatter that we observe is large and skewed towards values > 1 . This is in some part due to the effects of bandwidth- and time-smearing, which both reduce the peak flux densities of sources as a function of distance from the phase centre. Given the fact that the data were phase-shifted to the facet centres and averaged further, there may be additional smearing within the facet images. However, both the images and the time and frequency averaging are small enough that this should not be significant. Given the averaged channel, time and imaging resolution, we estimate (using the equations given by Bridle & Schwab 1989) the combined effect of bandwidth and time smearing due to the averaging in the full dataset to be of the order of 74 per cent (i.e. measured peak flux densities at 74 per cent of the expected value) at the edge of our field at 2.4° . Our empirical estimate for the correction factor for each source due to the combined effect is included in the source catalogue. While the traditional way of estimating this is to perform a Monte Carlo simulation in which fake sources are added to the uv -data and before calibration, this involves significant computational time and so has not been done here.

The effect of noise on the total-to-peak flux density ratios as a function of signal-to-noise can be determined by running complete simulations in which artificial sources are injected into the real data and imaged and detected in the same way as the observed data. However this requires re-running the full calibration and imaging routine, which we have not yet done due to computational constraints. Alternatively, it is common practice to determine a lower envelope, or flux-varying limit, above which a given fraction of sources lie and reflect it to obtain an upper envelope above which sources can be classified as resolved. In this case there is more scatter with $S_{\text{int}} > S_{\text{peak}}$, likely due to a further smearing effect which we have not accounted for. This may be a result of uncorrected residual phase errors away from the calibrator directions. We note though that the integrated flux densities are still correct – see Section 4.4.2. Thus, in preference to fitting a lower envelope, we have identified by eye a conservative upper envelope, defined by $S_{\text{int}}/S_{\text{peak}} > 1.25 + 250/SNR^{2.7}$, where $SNR = S_{\text{peak}}/\sigma$, and classified 1 158 sources above this envelope as resolved.

We have identified 187 large extended sources by visual inspection of the images. These sources span sizes (largest angular size, LAS) of ≈ 20 arcsec to ≈ 250 arcsec. The largest (LAS > 45 arcsec) of these sources are presented in Fig. 4.A.1 in the Appendix. In the cases of the largest sources these can appear as separate sources within our catalogue because the emission from the lobes of these giant radio galaxies is not contiguous.

Many diffuse extended sources are also clearly visible in the facet images. These sources can be missed as their peak flux densities fall below the detection threshold. We identified four very clear large diffuse sources – see Fig. 4.A.2 in the Appendix. A full study of diffuse emission is deferred to a subsequent paper, as this will require re-imaging of the facets with optimized parameters.

4.4.5 Source Catalogue

The final catalogue consists of 5 652 sources with flux densities between 0.4 mJy and 5 Jy and is available as part of the online version of this article. The astrometry in the catalogue has been corrected for the systematic offset described in Section 4.4.2. Both the integrated and peak flux densities in the catalogue have been corrected for the systematic offset (Section 4.4.2). Resolved sources are identified as described in Section 4.4.4. A sample of the catalogue is shown in Table 4.2 where the columns are:

Column (1) – IAU Source name,




Figure 4.16: The measured ratio of integrated to peak flux density as a function of signal-to-noise ratio. The dotted blue line shows the conservative upper envelope identified by eye above which sources are likely to be fully resolved.

Columns (2) and (3) – flux-weighted right ascension (RA) and uncertainty,
 Columns (4) and (5) – flux-weighted declination (DEC) and uncertainty,
 Column (6) – integrated source flux density and uncertainty,
 Column (7) – peak flux density and uncertainty,
 Column (8) – approximate correction factor to the peak flux density to account for bandwidth- and time-smearing,
 Column (9) – the local rms noise used for the source detection,
 Column (10) – number of Gaussian components,
 Column (11) – Flag indicating the resolved parameterisation of the source. ‘U’ refers to unresolved sources and ‘R’ to resolved sources.

4.5 Results

In this section, we report two results based on the LOFAR catalogue: the spectral indices between 150 and 1400 MHz, and the 150 MHz faint source counts. Further analysis of these data will be presented in future publications.

4.5.1 Spectral Index Distributions

We use the deep WSRT 1.4 GHz data covering the Boötes Field (de Vries et al. 2002) to calculate spectral indices between 150 MHz and 1.4 GHz, α_{150}^{1400} , the distribution of which is shown in Fig. 4.17. The WSRT map has a resolution of 13×27 arcsec, so some sources appear as separate sources in the LOFAR map but are identified as single sources in the WSRT image. To exclude erroneous spectral indices derived for such sources, we limit this selection to sources that are not

Table 4.2: Sample of the LOFAR 153 MHz source catalogue.

Source ID	RA [deg]	σ_{RA} [arcsec]	DEC [deg]	σ_{DEC} [arcsec]	S_{int} [mJy]	S_{peak} [mJy beam ⁻¹]	F_{smear} [mJy beam ⁻¹]	rms [mJy beam ⁻¹]	Gaussians	Resolved
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
J143506.04+363241.0	218.78	0.49	36.54	0.56	4.25 ± 0.75	1.74 ± 0.26	1.30	0.19	4	U
J143302.77+341853.1	218.26	2.20	34.31	1.08	4.25 ± 0.44	0.56 ± 0.11	1.00	0.12	1	R
J143635.58+334506.5	219.15	0.71	33.75	0.53	4.24 ± 0.49	1.04 ± 0.17	1.09	0.13	2	R
J143126.36+344109.8	217.86	1.51	34.69	1.36	4.24 ± 0.44	0.49 ± 0.11	1.00	0.12	1	R
J142244.06+354749.5	215.68	0.45	35.80	0.45	4.24 ± 0.78	2.75 ± 0.35	1.36	0.22	3	U
J143504.13+354743.8	218.77	0.45	35.80	0.47	4.24 ± 0.55	2.17 ± 0.26	1.13	0.14	3	R
J142730.62+334120.5	216.88	0.47	33.69	0.42	4.24 ± 0.47	2.80 ± 0.31	1.09	0.13	1	U
J143037.25+324208.2	217.66	1.20	32.70	1.73	4.24 ± 0.46	0.82 ± 0.18	1.21	0.17	1	U
J142737.16+344313.8	216.90	0.42	34.72	0.44	4.24 ± 0.51	3.13 ± 0.34	1.05	0.12	2	U
J143415.72+352456.8	218.57	0.53	35.42	0.49	4.24 ± 0.48	1.93 ± 0.25	1.06	0.15	1	R




Figure 4.17: Spectral index, α_{153}^{1400} , distribution of sources matched between 1.4 GHz and 153 MHz (points). The difference in resolution is 13×27 arcsec (WSRT) and 5.6×7.4 arcsec (LOFAR). The blue points show upper limits to the spectral index for LOFAR sources which do not have a higher frequency counterpart. The horizontal line shows the median spectral index of -0.79 determined using LOFAR sources with flux densities greater than 2 mJy.

identified as extended in either the LOFAR or WSRT catalogues and that do not have multiple matches within a 30 arcsec search radius.

Using LOFAR sources with flux densities greater than 2 mJy, we find a median spectral index between 1400 and 153 MHz of -0.79 ± 0.01 and scatter of $\sigma = 0.30$ which is consistent with previously reported values: -0.87 ± 0.01 , (Williams et al. 2013), -0.79 (Intema et al. 2011), -0.78 (Ishwara-Chandra et al. 2010), -0.82 (Sirothia et al. 2009), and -0.85 (Ishwara-Chandra & Marathe 2007). A detailed spectral-index analysis using the other available radio data is deferred to later works.

4.5.2 Source Counts

We used the LOFAR catalogue to compute the 150 MHz source counts down to ≈ 1 mJy. This is at least an order of magnitude deeper than previously studied at these low frequencies (e.g. McGilchrist et al. 1990; Intema et al. 2011; Williams et al. 2013). The source counts are computed using the integrated flux densities, but sources are detected based on their *measured* peak flux density over the local noise level. Thus, the completeness of the source counts depends both on the variation of the noise in the image and on the relation between integrated and peak flux densities. The latter is dependant both on systematic effects (e.g. smearing) and the intrinsic relation between integrated and peak flux densities of radio sources due to their intrinsic sizes. In the following paragraphs we discuss these effects and how we correct for them in deriving the source counts.

Visibility Area

Due to the large variation in rms across the single pointing image (see Fig. 4.11), sources of different flux densities are not uniformly detected across the image, i.e. faint sources can only be detected in a smaller area in the inner part of the image. Moreover, smearing causes a reduction in peak flux density while conserving the integrated flux density, and the amount of smearing depends on the distance to the phase-centre. We have noted already the effect of bandwidth- and time-smearing (see Section 4.4.4) and use the equations given by Bridle & Schwab (1989) to calculate an approximate correction to the peak flux density of each source based on its position in the map. The maximal correction is at most $S_{\text{peak}}^{\text{meas}} \approx 0.74 S_{\text{peak}}^{\text{corr}}$, so sources with *corrected* peak flux densities $> 6.7\sigma$ will have effective measured peak flux densities above the 5σ detection threshold. We therefore select only sources based on this threshold for deriving the source counts. To correct for the varying rms, we weight each source by the reciprocal of the area in which it can be detected, its visibility area, (e.g. Windhorst et al. 1985), based on its smearing-corrected peak flux density value. This also accounts for the varying detection area within a given flux-density bin.

Systematic Effects

Another effect that could potentially influence both the peak and integrated flux densities is CLEAN bias, which could bias both downwards at the lowest flux levels, thus leading to low source counts. However, we expect this to be negligible because the use of masks in the imaging and good uv -coverage should minimise the CLEAN bias. In general, noise can scatter sources into adjacent bins, again most noticeably at low flux densities. A positive bias is introduced by the enhancement of weak sources by random noise peaks (Eddington bias; Eddington 1913). Both of these effects could be quantified by simulations, but our source counts are not corrected for them, due to the computational expense of running the full required simulation.

Resolution Bias

A resolved source of a given integrated flux density will be missed by the peak-flux-density selection more easily than a point source of the same integrated flux density. This incompleteness is called the resolution bias and to make a correction for it requires some knowledge of the true angular size distribution of radio sources. We have estimated a correction for the resolution bias following Prandoni et al. (2001). First we calculate the approximate maximum size θ_{max} a source could have for a given integrated flux density before it drops below the peak-flux detection threshold. Fig. 4.18 shows the angular size of the LOFAR sources. We use the relation

$$\frac{S_{\text{int}}}{S_{\text{peak}}} = \frac{\theta_{\text{min}}\theta_{\text{max}}}{b_{\text{min}}b_{\text{max}}},$$

where b_{min} and b_{max} are the synthesized beam axes and θ_{min} and θ_{max} are the source sizes, to estimate the maximum size a source of a given integrated flux density can have before dropping below the peak-flux detection threshold. Given this θ_{max} we estimate the fraction of sources with angular sizes larger than this limit using the assumed true angular size distribution proposed by Windhorst et al. (1990): $h(> \theta_{\text{lim}}) = \exp[-\ln 2(\theta_{\text{lim}}/\theta_{\text{med}})^{0.62}]$ with $\theta_{\text{med}} = 2S_{1.4\text{GHz}}^{0.30}$ arcsec (with S in mJy; we have scaled the 1.4 GHz flux densities to 150 MHz with a spectral index of -0.8). We have also calculated the correction using $\theta_{\text{med}} = 2$ arcsec for sources with $S_{1.4\text{GHz}} <$




Figure 4.18: Angular size, θ (geometric mean), for the LOFAR sources as a function of integrated flux density. The blue shaded region shows the range of maximum size (θ_{\max}) a source of a given integrated flux density can have before dropping below the peak-flux detection threshold (the range reflects the range of rms noise in the LOFAR map). The red lines show the two functions used for the median angular size (θ_{med}) as a function of integrated flux density.

1 mJy (see Windhorst et al. 1993; Richards 2000). The resolution bias correction $c = 1/[1 - h(> \theta_{\max})]$ is plotted in Fig. 4.19 for the two different assumed distributions. In correcting the source counts we use an average of the two functions. We use the uncertainty in the forms of θ_{med} and in θ_{lim} to estimate the uncertainty in the resolution bias correction. We further include a overall 10 per cent uncertainty following Windhorst et al. (1990). While we have used the extrapolated Windhorst et al. (1990) size distribution from 1.4 GHz to correct the source counts presented here, we note that the observed size distribution (see Appendix 4.B) suggests that the low frequency emission is more extended. Thus the real resolution bias correction factor is likely to be somewhat larger, particularly in the lowest flux density bins, and may explain the turndown in source counts (see Fig. 4.20). A full study of the true low frequency angular size distribution of radio sources is beyond the scope of this paper.

Complex Sources

The source counts need to be corrected for multi-component sources, i.e. cases where the radio-lobes are detected as two separate sources. The flux densities of physically related components should be summed together, instead of counted as separate sources. We use the method described in White et al. (2012) and Magliocchetti et al. (1998) to identify the double and compact source populations. This is done by considering the separation of the nearest neighbour of each component and the summed flux of the component and its neighbour. Pairs of sources are regarded as single sources if the ratio of their flux densities is between 0.25 and 4, and their separation is




Figure 4.19: Resolution bias correction $1/[1 - h(> \theta_{\text{lim}})]$ for the fraction of sources with angular size larger than θ_{max} at a given integrated flux density. For the faintest sources two curves are shown: the blue curve shows θ_{med} as a function of S_{int} and the green curve shows the result assuming $\theta_{\text{med}} = 2$ arcsec at these flux densities (the range of each reflects the range of rms noise in the LOFAR map).

less than a critical value dependent on their total flux density, given by

$$\theta_{\text{crit}} = 100 \left[\frac{S}{10} \right]^{0.5},$$

where S is in mJy and θ is in arcsec. Approximately 460 sources in the sample used for calculating the source counts, or 8.5 per cent, are considered to be a part of double or multiple sources for the source count calculation.

The Low-Frequency mJy Source Counts

The Euclidean-normalized differential source counts are shown in Fig. 4.20. Uncertainties on the final normalised source counts are propagated from the errors on the correction factors and the Poisson errors (Gehrels 1986) on the raw counts per bin. Table 4.3 lists (1) the flux density bins, (2) the central flux density of the bin, (3) the raw counts, (4) the effective detection areas for sources at the lower and upper limits of the flux density bin where they are different, (5) the effective area corresponding to the bin centre, (6) the mean weight of the sources in the bin, (7) the resolution bias correction factor, and (8) the corrected normalised source counts.

Model source counts have been derived by Wilman et al. (2008) for the 151 MHz and 610 MHz source populations predicted from the extrapolated radio luminosity functions of different radio sources in a Λ CDM framework. We show the source counts for both AGN and star-forming (SF) galaxies on Fig. 4.20. The Wilman et al. (2008) model catalogue has been corrected with their recommended default post-processing, which effectively reduces the source count slightly at low flux densities. At low flux densities it is likely that the Wilman et al. (2008)




Figure 4.20: Euclidean-normalized differential source counts for the LOFAR 153 MHz catalogue (filled black circles) between 1 mJy and 5 Jy. Open black circles show the source counts without the resolution bias correction. For comparison we plot the SKADS model counts (Wilman et al. 2008) separated into AGN and SF for 151 MHz (in grey) and scaled from 610 MHz (in black). We show also the measured 153 MHz source counts (Williams et al. 2013) (in green) and scaled from 325 MHz (Mauch et al. 2013) (in red). Additionally, we have included several source counts determinations from 1.4 GHz from the compilation of de Zotti et al. (2010) scaled down to 150 MHz assuming a spectral index of -0.8 (grey points). The blue lines at the low flux density end show how the higher frequency counts would scale if a flatter spectral index of -0.5 was used, i.e. the tail of the source counts would lie further down and to the left in the plot. The vertical dashed lines show the minimum integrated flux density at which a point source could be detected in 10, 50 and 90 per cent of the map.

Table 4.3: Euclidean-normalized differential source counts for the LOFAR 150 MHz catalogue.

S Range [mJy]	S_c [mJy]	Raw Counts	Area [deg ²]	$A(S_c)$ [deg ²]	$\langle W \rangle$	Correction	Normalised counts [Jy ^{3/2} sr ⁻¹]
(1)	(2)	(3)	(4)	(4)	(5)	(6)	(7)
1.13 – 1.31	1.22	284 ⁺¹⁸	14.43 – 17.25	15.85	1.84	1.3 ± 0.3	32 ⁺¹
1.31 – 1.52	1.41	314 ⁺¹⁹	17.25 – 18.80	18.33	1.49	1.2 ± 0.2	34 ⁺¹
1.52 – 1.78	1.65	365 ⁺²⁰	18.80 – 19.17	19.06	1.54	1.2 ± 0.2	47 ⁺¹
1.78 – 2.11	1.94	374 ⁺²⁰	19.17 – 19.24	19.22	1.29	1.2 ± 0.2	47 ⁺²
2.11 – 2.53	2.31	385 ⁺²¹	19.24 – 19.27	19.26	1.19	1.1 ± 0.1	53 ⁺²
2.53 – 3.08	2.80	369 ⁺²⁰	19.27 – 19.29	19.28	1.11	1.1 ± 0.1	57 ⁺³
3.08 – 3.81	3.43	298 ⁺¹⁸	19.29 – 19.30	19.29	1.07	1.1 ± 0.1	54 ⁺³
3.81 – 4.79	4.27	299 ⁺¹⁸	...	19.30	1.08	1.1 ± 0.1	71 ⁺⁴
4.79 – 6.16	5.44	230 ⁺¹⁶	...	19.30	1.02	1.1 ± 0.1	65 ⁺⁴
6.16 – 8.13	7.08	220 ⁺¹⁶	...	19.30	1.04	1.0 ± 0.1	87 ⁺⁶
8.13 – 11.05	9.48	218 ⁺¹⁶	...	19.30	1.03	1.0 ± 0.1	116 ⁺⁸
11.05 – 15.59	13.13	177 ⁺¹⁴	...	19.30	1.01	1.0 ± 0.1	133 ⁺¹¹
15.59 – 23.04	18.95	169 ⁺¹⁴	...	19.30	1.00	1.0 ± 0.1	192 ⁺¹⁶
23.04 – 36.09	28.83	165 ⁺¹⁴	...	19.30	1.00	1.0 ± 0.1	305 ⁺²⁶
36.09 – 60.94	46.89	143 ⁺¹³	...	19.30	1.00	1.0 ± 0.1	470 ⁺⁴²
60.94 – 113.67	83.23	122 ⁺¹²	...	19.30	1.00	1.0 ± 0.1	786 ⁺⁷⁸
113.67 – 242.77	166.12	83 ⁺¹⁰	...	19.30	1.00	1.0 ± 0.1	1233 ⁺¹⁵⁰
242.77 – 626.60	390.03	49 ⁺⁸	...	19.30	1.00	1.0 ± 0.1	2062 ⁺³³⁹
626.60 – 2124.44	1153.76	20 ⁺⁶	...	19.30	1.00	1.0 ± 0.1	3247 ⁺⁹⁰²

counts slightly overestimate the true counts due to double counting of hybrid AGN-SF galaxies. These models are based on low-frequency data at higher flux density limits and higher frequency data so some deviations are not unexpected; however, our observed counts do follow their model quite well. Mauch et al. (2013) suggest that the spectral curvature term used in the Wilman et al. (2008) models mean that their 151 MHz counts under-predict reality. For this reason we include the model counts at both frequencies.

Source counts below 1 mJy at 1.4 GHz have been the subject of much debate. For comparison, in Fig. 4.20, we have included the several source count determinations from 1.4 GHz scaled down to 150 MHz from the compilation of de Zotti et al. (2010) (including counts from Bridle et al. 1972; White et al. 1997; Ciliegi et al. 1999; Gruppioni et al. 1999; Richards 2000; Hopkins et al. 2003; Fomalont et al. 2006; Bondi et al. 2008; Kellermann et al. 2008; Owen & Morrison 2008; Seymour et al. 2008). This is a representative comparison and not an exhaustive list of available source counts. In particular, there are even deeper models of higher frequency counts using statistical methods (e.g. Vernstrom et al. 2014; Zwart et al. 2015). The source counts are scaled assuming a spectral index of -0.8 . The blue lines at the low flux density end show how the higher frequency counts would scale if a flatter spectral index of -0.5 was used, i.e. the blue line drawn through the tail of the source counts would lie at lower 147 MHz fluxes (left in the plot) and at lower normalised count values (down in the plot) mostly due to the $S^{5/2}$ term in the normalised counts. We note that the flattening of the source counts at 5 – 7 mJy, associated with the growing population of SF galaxies and faint radio-quiet AGN at lower flux densities (see e.g. Jarvis & Rawlings 2004; Simpson et al. 2006; Padovani et al. 2015), is clear and is the same as that seen at the higher frequencies. The further drop in the lowest flux density bins may be the result of some unaccounted for incompleteness in our sample or different resolution bias correction (see Section 4.5.2 and Appendix 4.B).

4.6 Conclusion

We have presented LOFAR High Band Antenna observations of the Boötes field made as part of the LOFAR Surveys Key Science Project. These are the first wide area (covering 19 deg^2), deep (reaching $\approx 120 - 150 \mu\text{Jy beam}^{-1}$), high resolution ($5.6 \times 7.4 \text{ arcsec}$) images of one of the extragalactic deep fields made at 130 – 169 MHz. These observations are at least an order of magnitude deeper and 3 – 5 times higher in resolution than previously obtained at these frequencies. We have used a new calibration and imaging method to correct for the corrupting effects of the ionosphere and LOFAR digital beams.

The radio source catalogue presented here contains 5 652 sources detected with peak flux densities exceeding 5σ . We have quantified the positional and flux density accuracy of the LOFAR sources and used the source catalogue to derive spectral indices between 150 and 1400 MHz, finding a median spectral index of -0.79 ± 0.01 . Finally, we have presented the deepest differential source counts at these low frequencies. These source counts follow quite well the model predictions of Wilman et al. (2008) and show the flattening at a few mJy as a result of the increasing contribution of SF galaxies.

Acknowledgements

LOFAR, the Low Frequency Array designed and constructed by ASTRON, has facilities in several countries, that are owned by various parties (each with their own funding sources), and

that are collectively operated by the International LOFAR Telescope (ILT) foundation under a joint scientific policy. The Open University is incorporated by Royal Charter (RC 000391), an exempt charity in England & Wales and a charity registered in Scotland (SC 038302). The Open University is authorized and regulated by the Financial Conduct Authority.

Appendix

4.A Postage stamp images

The largest ($LAS > 45$ arcsec) of the visually identified extended sources are shown in Fig. 4.A.1, and the four very clear large diffuse sources are shown in Fig. 4.A.2.

4.B Source size distribution

We have investigated the extent to which the extrapolation of the Windhorst et al. (1990) size distribution might be valid at low frequencies. This is relevant to the resolution bias correction to the source counts, described in Section 4.5.2. We have done this by comparing the true angular size distribution of the 147 MHz LOFAR sources to the Windhorst et al. (1990) distribution, given by

$$h(> \theta_{\text{lim}}) = \exp[-\ln 2(\theta_{\text{lim}}/\theta_{\text{med}})^{0.62}],$$

with the median size, in arcsec, as a function of flux density of

$$\theta_{\text{med}} = 2S_{1.4\text{GHz}}^{0.30},$$

where S is in mJy and we have scaled the 1.4 GHz flux densities to 150 MHz with a spectral index of -0.8 . The observed and extrapolated size distributions are shown in Fig. 4.B.1 for four flux density bins. The low frequency emission appears to be more extended, which would suggest that the actual resolution bias correction should be somewhat larger. Future LOFAR Survey results, in particular, at high resolution using very long baseline interferometry with the LOFAR international stations, will allow for more detailed studies of the true angular size distribution of radio sources.




Figure 4.A.1: Postage stamps of extended sources identified visually, showing all the sources with approximate $LAS > 45\text{arcsec}$. These include some Giant Radio Galaxies. The greyscale shows the flux density from $-3\sigma_{\text{local}}$ to $30\sigma_{\text{local}}$ where σ_{local} is the local rms noise. The scalebar in each image is 1arcmin.




Figure 4.A.1: Continued.




Figure 4.A.1: Continued.




Figure 4.A.1: Continued.




Figure 4.A.2: Postage stamps of large diffuse sources identified by eye. The greyscale shows the flux density from $-3\sigma_{\text{local}}$ to $15\sigma_{\text{local}}$ where σ_{local} is the local *rms* noise. The scalebar in each image is 1 arcmin.




Figure 4.B.1: The observed source size distribution (in blue) in four flux density intervals. The magenta curves show the Windhorst et al. (1990) size distribution for the upper and lower bounds of the flux density bin. In each the vertical dotted lines are b_{maj} (on the left) and the approximate maximum size a source can have before it drops below the peak-flux detection threshold (on the right). The catalogue will be incomplete for sources larger than the right line.

