

The absolute flux scale at 142 MHz

Jeffrey Bout

Master Thesis

Supervisors: Dr. M. A. Brentjens & Prof. dr. A. G. de Bruyn

Kapteyn Astronomical Institute, University of Groningen

January 22, 2010

Abstract

The absolute flux scale below 408 MHz is currently rather uncertain, with typical uncertainties between 300 and 20 MHz in the range of 10-20%. Accurate spectral index work requires uncertainties well below 5%. The uncertainties can be reduced by adding new observations at low frequencies of radio sources acting as calibrators. With the arrival of digitized processing and calibration techniques, precise flux ratio measurements on bright radio sources should be feasible, improving the spectra of calibrators at low frequencies. Using WSRT/LFFE, we have determined flux ratios of seven secondary calibrators. These flux ratios could not be transformed to absolute flux ratios.

Contents

1	Introduction	5
1.1	Obtaining calibrators in the 100 MHz regime	6
1.2	The WSRT and the LFFE	7
2	Observations	11
2.1	Properties of individual observations	11
2.2	Observing sessions	13
3	Pre-processing	17
3.1	Hanning filtering	18
3.2	System temperature corrections	20
3.3	Data quality	23
3.4	WSRT flagger	28
3.5	Additional flagging	29
3.6	Scale visibility amplitudes per field	31
3.7	Renaming moved antennas	32
4	Self calibration	35
4.1	Formulating the measurement equation for the WSRT	36
4.2	Measurement equation for our observations	38
4.3	Local Sky Models for our fields	41
4.4	Setting up self calibration in MeqTrees	46
4.5	Step 1: Bandpass calibration	47
4.6	Step 2: Phase calibration	50
4.7	Step 3: Flux calibration	52
4.8	Step 4: Repeat phase and flux calibration	54
4.9	Step 5: Gain calibration	55
4.10	Self calibration quality	58
5	Data analysis	63
5.1	Obtaining relative fluxes	63
6	Conclusions and suggestions for future work	67

7	Acknowledgments	69
A		71
A.1	Scripts	71
A.2	Color palette for displaying complex values	72
A.3	Figures	73

Chapter 1

Introduction

After many years of absence, radio astronomy at low frequencies now has an increasing popularity with the construction of new instruments like the Low Frequency Array (LOFAR) which operates at 15 to 240 MHz (Rottgering et al., 2006). Observing at low frequencies brings along the need for a precise flux scale at these frequencies. The uncertainties in flux scales at low frequencies is still big: in the range of 10 to 20% between 20 and 300 MHz. For proper operation of instruments like LOFAR, this uncertainty should be reduced to less than 5%.

The flux scale at low frequencies is presently determined by extrapolating high frequency spectra (above 408 MHz) of several bright radio sources combined with flux ratio measurements at a few lower frequencies. Each frequency was observed with a different instrument. The resulting flux scale becomes less accurate with lower frequencies: the flux scales of the Very Large Array (VLA) and the Westerbork Synthesis Radio Telescope (WSRT) already disagree by 7% at 327 MHz.

One way to reduce the uncertainties in the low frequency flux scales is to increase the number of observations in this regime on radio sources acting as calibrators. The observations on which the current flux scales are based were done in the sixties (Kellermann et al., 1969) and seventies (Baars et al., 1977). Over the last two decades big improvements have been achieved in radio astronomy with the arrival of digitized processing and calibration techniques. Therefore more precise flux measurements on bright radio sources should be feasible, improving their spectra at low frequencies.

This report aims at delivering improved relative fluxes of calibrators at low frequencies. It describes relative flux measurements using observations from 2006 on 13 radio sources. The observations were done at 115 to 180 MHz with the Low Frequency Front Ends (LFFE) on the WSRT. The observations were proposed by A. G. de Bruyn and R. G. Strom with project number DEV/203. The observations are described in the next chapter.

Before the relative fluxes are measured the observations are corrected in two phases. The observations are pre-processed (for example unreliable data is removed) and self calibrated afterwards. The two phases are described in the next two chapters.

Table 1.1: *Properties of the observed sources. Right ascension and declination are in ICRS coordinates (epoch=2000.0). The spectral index curvature represents the deviation from a straight line of the absolute flux density in a log-log-plot. Sources with increasing spectral indices at low frequencies show negative curvatures in the log-log-plots of their absolute flux densities. Sources of which their spectral indices change sign at low frequencies show a cutoff in the log-log-plots of their absolute flux densities. Spectral index curvatures for Cyg A and Tau A are from Baars et al. (1977) and from Kellermann et al. (1969) for all the other sources. All sources without spectral index curvature at low frequencies are used as primary calibrators, others as secondary calibrators.*

Source	Right Ascension	Declination	Type	Spectral index curvature
3C48	01 37 41.30	+33 09 35.1	Quasar	Cutoff
3C123	04 37 04.37	+29 40 13.8	Radio galaxy	Negative
3C147	05 42 36.13	+49 51 07.2	Seyfert 1	Cutoff
3C196	08 13 36.03	+48 13 02.5	Seyfert 1	Negative
3C295	14 11 20.64	+52 12 09.1	Seyfert 1	Cutoff
3C348	16 51 09.00	+04 59 34.0	Radio galaxy	None
3C353	17 20 28.16	-00 58 46.6	Radio galaxy	None
3C380	18 29 31.78	+48 44 46.1	Seyfert 1	Positive
3C409	20 14 27.74	+23 34 58.4	Unidentified	Negative
Cas A	23 23 24.00	+58 48 54.0	SNR	None
Cyg A	19 59 28.35	+40 44 02.0	Seyfert 2	Negative
Tau A	05 34 31.97	+22 00 52.1	SNR / Pulsar	None
Vir A	12 30 49.42	+12 23 28.0	AGN	None

The corrected observations are analyzed in Chapter 5. The relative fluxes of seven sources have been measured. These relative fluxes could not be transformed to absolute flux ratios.

1.1 Obtaining calibrators in the 100 MHz regime

We want to obtain absolute flux densities of a number of radio sources. Future instruments like LOFAR can then use these fluxes for calibration. Unfortunately we do not have the ability to directly measure absolute flux densities. The necessary model for the absolute gains of WSRT/LFFE is not available.

Absolute flux densities are obtained indirectly by first measuring relative fluxes of a number of sources and then comparing these fluxes with absolute flux densities from the literature. If a subset of sources are fit to the literature values using one scaling factor, the scaling factor for transforming all the unitless fluxes to absolute flux densities is found.

To use the indirect method at least a number of sources should be observed with well known absolute flux densities at our frequency. And indeed these absolute flux densities

can be obtained because in radio astronomy some sources are known that have a spectrum behaving like a power-law:

$$\text{Flux}(\nu) \propto \nu^\alpha, \quad (1.1)$$

where α is called the *spectral index* of the source. The radio emission of these sources is typically generated by the synchrotron mechanism, in which high-energy electrons in magnetic fields radiate as a result of their orbital motion. If the spectral index is constant over a frequency range, the absolute flux density at any frequency within that range becomes very predictable. The log-log-plot of the absolute flux density will show a straight line with slope α . Some sources, however, will begin to deviate at low frequencies or even show a cutoff in the log-log-plot. This deviation is caused by synchrotron self-absorption, and is produced at low frequencies if the electron density and magnetic fields are high enough. Therefore the absolute flux density becomes less predictable at low frequencies.

Radio sources 3C348, 3C353, Cas A, Tau A, and Vir A are known to have constant spectral indices at low frequencies (Kellermann et al., 1969; Baars et al., 1977). They are the primary calibrators within our set of observed sources.

For calibrating instruments like LOFAR one needs more than just five calibrators. Eight secondary calibrator candidates were selected: 3C48, 3C123, 3C147, 3C196, 3C295, 3C380, 3C409, and Cyg A, adding up to 13 sources. The properties of all the sources are given in Tab. 1.1. The eight new sources do not have constant spectral indices, but are expected to be strong in order to dominate their fields-of-view. This way the measured fluxes have little interference from surrounding sources. Maps of some of the fields-of-view at different frequencies are shown in Fig. 4.4 and Fig. A.3. Cyg A is a very important strong source. According to Kellermann et al. (1969) it is the second strongest source in the sky at both 38 MHz and 178 MHz. 3C 147, 3C 48 and 3C 295 are well known calibrators at higher frequencies. Finally, 3C196 and 3C295 are interesting because they are far away from the galactic plane, in the north galactic halo.

1.2 The WSRT and the LFFE

Our observations were conducted with the Westerbork Synthesis Radio Telescope (WSRT). This telescope is an aperture synthesis interferometer that consists of a linear array of 14 identical antennas with 25 m dishes arranged on a 2.7 km east-west line. The antennas are equipped with Multi Frequency Front Ends (MFFE) covering eight frequency bands between 240 MHz and 9 GHz. Since 2004 the 14 antennas are also equipped with Low Frequency Front Ends (LFFE) covering the frequency band from 115 MHz up to 180 MHz. Ten of the antennas are on fixed mountings and four antennas are movable along two railtracks. The 10 antennas on fixed mountings (labeled 0 to 9) are 144 m apart. The four remaining antennas (labeled A, B, C, D) are movable along two railtracks. One, 300 m long, is next to the fixed antennas and the other, 180 m long, is 9 x 144 m eastwards. The baselines can extend from 36 m to 2.7 km.

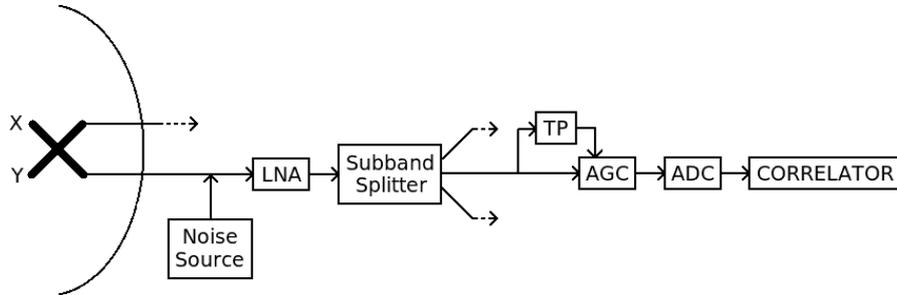


Figure 1.1: *Block diagram of one polarization and one subband of the WSRT hardware. LNA = Low Noise Amplifier, TP = Total Power Measurement, AGC = Automatic Gain Control, ADC = Analog to Digital Converter. Noise injection by a noise source is discussed in Sect. 3.2*



Figure 1.2: *An LFFE consists of four folded dipole antennas: two for each linear polarization. The dipoles can be moved to prevent blocking the MFFE (photo Astron).*

The WSRT antennas are equatorially mounted. This type of mounting ensures fixed orientations of the receiving dipoles with respect to the sky.

The Digital Backend (DZB, Z = last) correlator of the WSRT is a digital XF correlator. Its output is divided into eight frequency bands, each with a minimum of 64 channels. The bandwidth of each subband depends on the used front end. The correlator can deliver up to four polarizations for doing full linear polarimetry.

A block diagram of one polarization and subband of the WSRT hardware is shown in Fig. 1.1. After the incoming signal is amplified by the LNA and split into eight subbands, the total power is measured. With the measured power the signal is stabilized in the Automatic Gain Control (AGC) and digitized by the ADC. The digital signal is fed to the correlator. Stabilization of the signal by the AGC is discussed in detail in Sect. 3.2.

Since 2004 the 14 antennas of the WSRT are equipped with Low Frequency Front Ends (LFFE) covering the frequency band from 115 MHz up to 180 MHz. The boundaries are defined on the lower side by the FM band and on the upper side by a strong local TV transmitter (182 -188 MHz). The TV transmitter has been switched off since 30 October 2006, too late for our observations. An LFFE consists of four folded dipole antennas: two for each linear polarization.

The LFFEs were installed on the WSRT to extend the frequency coverage towards low frequencies. The LFFE operates in a part of the high frequency band of the Low Frequency Array (LOFAR). The LFFEs were developed to gain experience in the LOFAR frequency regime.

The frequency regime of the LFFEs is heavily polluted by RFI. To avoid RFI, the eight subbands can be positioned on RFI poor parts of the frequency band. But still RFI appears in the selected subbands. However, this RFI usually is narrow band and temporary. Therefore flagging this RFI is necessary. Also integration times of 10 seconds are recommended to limit the loss of data due to temporary RFI (and to detect fast phase shifts caused by the ionosphere, see Chapter. 4).

Chapter 2

Observations

2.1 Properties of individual observations

When observing fluxes of sources it is necessary to observe them under the same conditions. Following the typical observing strategy in radio astronomy of observing the fields in successive 12 hour exposures is unwise. With time many interfering effects vary like instrumental effects, atmospheric effects, and RFI. It is better to do short observations and switch between the fields as much as possible. This way successive targets will be observed with similar conditions. A drawback of switching between fields is the obtained uv coverage of each field. The image quality is imperfect when not all hour angels are used. Each of the observations took 15 minutes. Making the integration times even shorter would have lowered the total integration time. With a slew rate of 18° per minute on each axis it takes 9 minutes at maximum to switch between fields (the biggest movement on one axis made during our observations was 154°).

The integration times of the samples were 10 seconds, the shortest values possible for the WSRT. They were kept small to limit the loss of data due to intermittent RFI. Short integration times are also important because observations at low frequencies are affected

Table 2.1: *Central frequencies of the eight subbands.*

Subband ID	Central frequency (MHz)
0	117.0
1	130.0
2	139.5
3	142.0
4	147.5
5	156.0
6	162.0
7	174.0

by the ionosphere substantially. With short integration times, decorrelation due to rapid ionospheric phase fluctuations is greatly reduced.

Another argument for short integration times is fringe smearing in time. With short integration times the field-of-view is kept large. We will see that the contribution of bright sources at large angular distances is substantial at low frequencies (see Sect. 4.9 and Sect. 4.10). Therefore, by keeping the integration times short, we keep track of these bright sources with all baselines. For an integration time of 10 seconds, the maximum angular distance detected by all baselines is 57° , as shown below.

The pathlength difference of a baseline is:

$$l = d \sin \theta. \quad (2.1)$$

where d is the antenna distance in wavelengths and θ is the angle from the plane perpendicular to the baseline. The change of pathlength difference is therefore:

$$\frac{dl}{dt} = d \cos \theta \frac{d\theta}{dt}. \quad (2.2)$$

The oscillation time is equivalent to one over the change of pathlength difference. For an integration time of 10 seconds the shortest oscillation time that can be identified is 20 seconds, according to the Nyquist criterion. The value of $d\theta/dt$ can be as high as $7.27 \cdot 10^{-5}$ rad/s for objects on the celestial equator. Inserting this value into Eq. (2.2), together with an oscillation time of 20 seconds and the maximum WSRT antenna distance of 2.7 km or 1278 wavelengths at 142 MHz, θ becomes 57° . The frequency 142 MHz is the one selected for processing in Sect. 3.3.

The eight subbands of the WSRT backend were placed in the default areas with little RFI. The central frequencies of all the subbands are given in Tab. 2.1. The bandwidth of each subband was 1.25 MHz separated over 128 channels with widths of 9.8 KHz. Therefore only a small part (10 MHz) of the total bandwidth of the LFFE system (65 MHz) was covered. The channel width was chosen small to get as many RFI free channels as possible.

During observing only the xx and yy correlation products were stored. The xy and yx correlation products are less important if the observed sources are expected to be unpolarized and if the instrumental polarization is negligible. Unfortunately the latter condition is probably not met in our observations. This is discussed in detail in Sect. 4.2. The two correlation products were discarded in order to process 128 channels. Since 2006 the WSRT backend has been improved and nowadays recording all four polarizations is possible with even better spectral resolution.

For the observations all the other settings were typical: sources were tracked while observing and corresponding fringe stopping was applied. Total power measurements were done (more on this in Sect. 3.2) and phases were calibrated before the beginning of the observations.

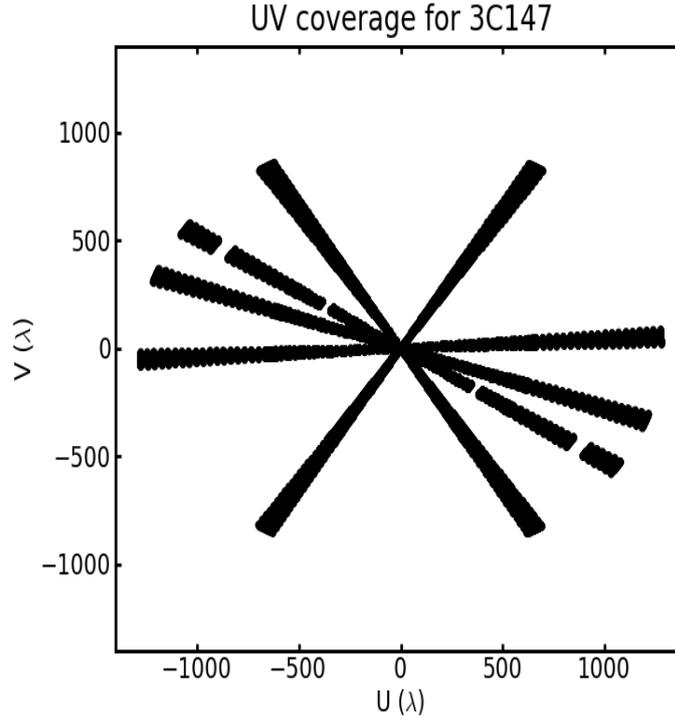


Figure 2.1: *The uv coverage at 142 MHz of 3C147 for the four sessions with usable data. The frequency 142 MHz is the one selected for processing in Sect. 3.3. The uv coverage of the other fields is shown in Fig. A.3*

2.2 Observing sessions

The observations were conducted during eight sessions spanning two and a half months in 2006. Each session took between 6.5 and 13 hours adding up to 70 hours in total. Properties of the sessions are given in Tab. 2.2. Some of the sessions were observed during daytime, hoping for little interference by the Sun. 2006 was in the solar minimum during which RFI from solar flares is less present.

Between the sessions, antennas A, B, C, and D were moved once. This affects the uv coverage and the data reduction, which is treated in Sect. 3.7.

Only four of the eight sessions contain useful data. The first, seventh, and eighth sessions were distorted by instrumental problems. The fifth session suffered from solar interference. Therefore the number of useful observations was much lower than originally hoped for. The uv coverage was correspondingly worse, leading to a bad point spread function (PSF) in the image plane. This can be problematic when locating sources during flux calibration in Sect. 4.7. Table 2.3 lists the number of successful observations per field. The uv coverage after excluding the four bad sessions is shown in Fig. 2.1 and Fig. A.3. The uv coverage is

particularly bad for fields 3C48, 3C348, and 3C353.

Table 2.2: *The properties of the eight observing sessions.*

Number	Starting time	Duration	Number of observations	Observed fields
1	24-07 23:02	6 ^h 55 ^m	21	3C123, 3C147, 3C353 3C380, 3C409, Cas A Cyg A, Tau A
2	26-07 23:34	6 ^h 34 ^m	20	3C48, 3C123, 3C147 3C353, 3C380, 3C409 Cas A, Cyg A, Tau A
3	30-07 22:58	6 ^h 55 ^m	21	3C48, 3C123, 3C147 3C353, 3C380, 3C409 Cas A, Cyg A, Tau A
4	06-08 00:28	9 ^h 15 ^m	27	3C48, 3C123, 3C147 3C196, 3C409, Cas A Cyg A, Tau A, Vir A
5	13-08 04:27	13 ^h 05 ^m	40	3C48, 3C123, 3C147, 3C196 3C295, 3C348, 3C353, 3C380 Cas A, Tau A, Vir A
6	02-09 08:58	9 ^h 12 ^m	28	3C123, 3C147, 3C196, 3C295 3C348, 3C353, 3C380 Cyg A, Tau A, Vir A
7	20-09 20:50	9 ^h 05 ^m	27	3C48, 3C123, 3C147, 3C196 3C353, 3C380, 3C409 Cas A, Cyg A, Tau A
8	11-10 15:03	9 ^h 02 ^m	27	3C123, 3C147, 3C295 3C348, 3C353, 3C380 3C409, Cas A, Cyg A

Table 2.3: *The number of observations for each field after dropping the four bad sessions in Sect. 3.3.*

Field	Number of Observations
3C48	3
3C123	14
3C147	7
3C196	6
3C295	6
3C348	4
3C353	5
3C380	7
3C409	14
Cas A	9
Cyg A	6
Tau A	6
Vir A	6

Chapter 3

Pre-processing

The raw observations produced by the WSRT contain many imperfections such as interference, gain variations, position- and time-variable system temperatures, non-uniform bandpasses, and ionospheric phase fluctuations. These need to be dealt with before one can extract reliable flux ratios from the data sets. The required data processing consists of two phases: pre-processing and self calibration.

The pre-processing phase is used to correct the data for known problems, apply external calibration information, and remove unreliable data. The purpose of the self calibration phase (Pearson and Readhead, 1984) is to derive corrections for instrumental and atmospheric effects based on a model of the sky, a model of the instrumental and atmospheric effects, and the observed Stokes parameters (Hamaker et al., 1996; Sault et al., 1996; Hamaker, 2000). Self calibration is covered in Chapter 4.

We start the pre-processing with two necessary steps that could be done while generating the observation data as well, but were omitted to do offline: tapering and inversion of automatic gain control (AGC) corrections.

Tapering is necessary to reduce the enormous spectral side lobes that exist in the lag spectra. They are a result of not tapering the correlated data before Fourier transforming the data by the XF correlator of the WSRT.

In the WSRT backend gain corrections are done to stabilize the antenna output powers, which are highly variable. The stabilization is necessary to make digitization independent of these variabilities in power. Inversion of the gain corrections is necessary to make flux measurements possible, which is the main goal of this project.

The next important step in pre-processing is remove unreliable data. Removing unreliable data is done after tapering and inversion of the gain corrections, because the spectrum is cleaner and the time series are smoother after these operations. Removing unreliable data is done in three phases: removing data on a high level (observing sessions and subbands), flagging corrupted visibilities using the WSRT flagger algorithm (Renting and Brentjens, 2006), and adding low level manual flags.

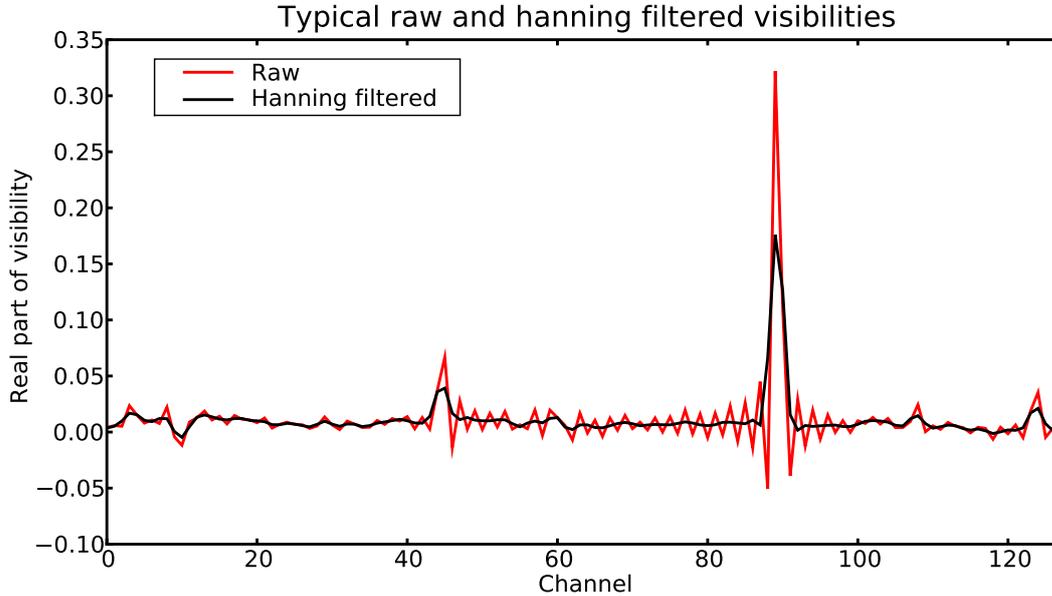


Figure 3.1: *Typical real parts of the visibilities as a function of frequency for one time slot and one subband. The red curve shows the raw data. Because no tapering was applied, sinc shaped patterns are visible. The black curve shows the data after tapering in the frequency domain using the Hanning filter. The side lobes have more or less disappeared, but the peaks have broadened.*

Before finishing the pre-processing, two steps are done to prepare our data for self calibration. The visibility amplitudes for each field are scaled so fitting a local sky model (LSM) to our data during self calibration is simplified. Also, the antenna names in our data are partly renamed. Antennas A-D are moved between two sessions and this is not expected by the used calibration package MeqTrees (Noordam and Smirnov, 2009).

3.1 Hanning filtering

The WSRT has a lag or XF correlator. It is a spectral correlator that measures the cross-correlation of two signals $V_1(t)$ and $V_2(t)$ as a function of time lag τ (Thompson et al., 2001):

$$V_1(t) \star V_2(t) = r(\tau) = \sum_{t=0}^{N-1} V_1(t) V_2^*(t - \tau), \quad (3.1)$$

where \star is the cross-correlation operator and N is the number of discrete points in the cross-correlation function. Fourier transforming the cross-correlation product results in the cross power spectrum. According to the convolution theorem:

$$V_1(t) \star V_2(t) \Rightarrow \widehat{V}_1(\nu) \widehat{V}_2^*(\nu), \quad (3.2)$$

where $\hat{\Rightarrow}$ represents Fourier transformation and $\hat{V}_1(\nu)$ is the Fourier transform of $V_1(t)$. Measurement of the cross-correlation is done over a limited time lag range, which is equivalent to applying a top-hat window of width $2N\tau_s$ (where τ_s is the sampling interval). This corresponds to convolving the cross power spectrum with the Fourier transform of the top-hat function, that is, with the sinc function:

$$\text{sinc}(\nu) = \frac{\sin(\pi\nu N/\Delta\nu)}{\pi\nu}, \quad (3.3)$$

which is normalized to unit area with respect to ν . The full width at half maximum (FWHM) of the sinc function is $1.21\Delta\nu/N$. Features within the spectrum are convolved with Eq. (3.3). Figure 3.1 shows an example of this. Bright RFI peaks have strong side lobes in the power spectrum and this is unwanted. Reducing the side lobes can be done by tapering the limited time lag cross-correlations, that is, by replacing the mentioned top-hat window function with a function with smaller side lobes in its Fourier transform. A well known tapering function is the Hanning window (Harris, 1978):

$$\omega(n) = \sin^2\left(\frac{n\pi}{N}\right) = 0.5 \left(1.0 - \cos\left(\frac{2n\pi}{N}\right)\right), n = 0, 1, \dots, N - 1. \quad (3.4)$$

Using this function reduces the first side lobe by a factor of 9 and degrades the resolution by 1.67, increasing the half-maximum level from $1.21\Delta\nu/N$ to $2.00\Delta\nu/N$.

Hanning tapering can be done both in the time domain and in the frequency domain. In practice, in the frequency domain the cross power spectrum should be convolved with the DFT of the Hanning window. This DFT is (Harris, 1978):

$$W(\theta) = 0.5D(\theta) - 0.25 \left[D\left(\theta - \frac{2\pi}{N}\right) + D\left(\theta + \frac{2\pi}{N}\right) \right], \quad (3.5)$$

where

$$D(\theta) = \exp\left(+j\frac{\theta}{2}\right) \frac{\sin\left(\frac{N}{2}\theta\right)}{\sin\left(\frac{1}{2}\theta\right)}. \quad (3.6)$$

The Hanning window is sampled at multiples of $2\pi/N$ and of these samples only three result in a non-zero Hanning DFT: $-2\pi/N$, 0 , and $+2\pi/N$. They are: $\frac{1}{4}$, $\frac{1}{2}$, and $\frac{1}{4}$. Therefore the DFT of the Hanning window is a three point triangular function and Hanning tapering before the DFT is equivalent to convolving the cross power spectrum with this simple triangular function after the DFT.

No tapering was done before doing the FFT for our observations. Therefore the raw data still has sinc shaped patterns as shown in Fig. 3.1 as the red curve. The black curve shows the result of Hanning tapering in the frequency domain. The side lobes more or less disappeared, but the peaks have broadened as well.

Flagging data (such as the big RFI peak in Fig. 3.1) should be done after tapering. If done before tapering, the side lobes could have been flagged while they more or less disappear during tapering.

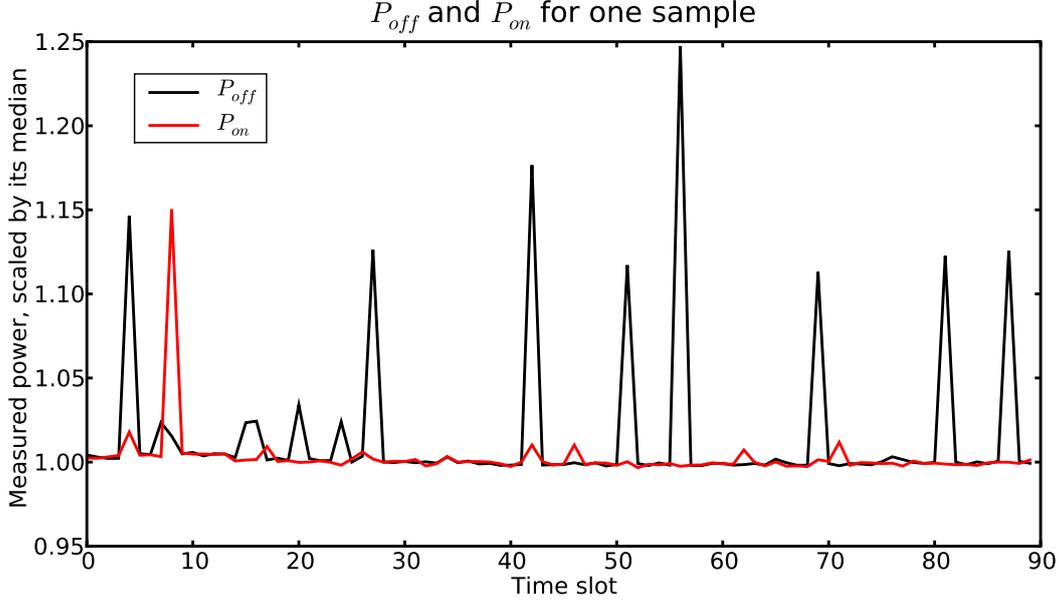


Figure 3.2: Successive P_{off} (black curve) and P_{on} (red curve) measurements of a typical observation for one antenna, subband, and polarization. Both curves are scaled by their median values. The peaks in both curves are caused by RFI and this RFI seems mostly uncorrelated.

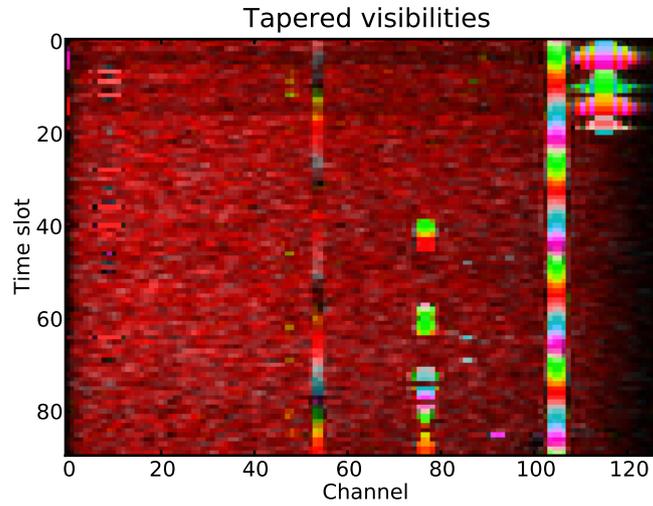
3.2 System temperature corrections

ADCs digitize the signal entering the WSRT for each antenna, polarization, and subband. This is shown in the block diagram in Fig. 1.1. The ADCs have 1.5 bit sampling (smaller than $-a$, between $-a$ and $+a$, larger than $+a$, where a is a fixed input voltage). However, the signal power has variations caused by the different powers of astronomical sources, time-varying RFI, and time-varying instrumental and atmospheric effects. To stabilize the input of ADCs, the signal powers are measured and fed into automatic gain control (AGC) units installed in front of the ADCs. The gains applied by the AGCs are stored together with the visibilities. Therefore deriving the original powers can be done off-line by inverting the gain corrections of the AGCs.

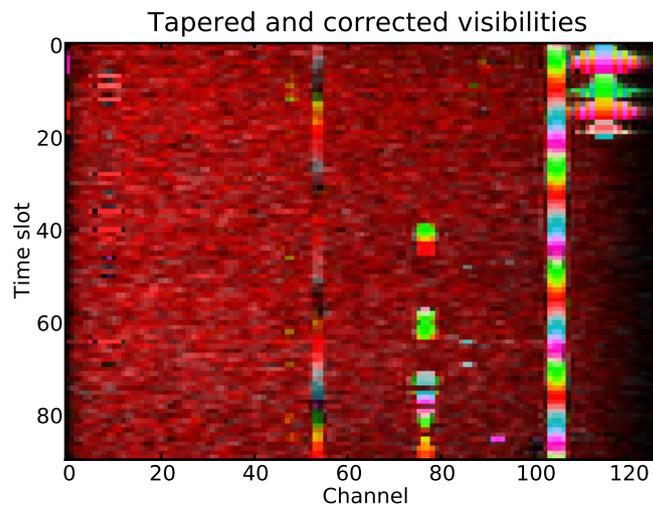
The measured powers are unitless numbers. To determine the absolute powers (or their corresponding system temperatures T_{sys}), the signals of noisesources with known power (P_{noise}) are periodically added to the astronomical signals. The system temperature can then be derived, using:

$$P_{\text{on}} = P_{\text{off}} + P_{\text{noise}}, \quad (3.7)$$

$$\frac{T_{\text{sys}}}{T_{\text{noise}}} = \frac{P_{\text{off}}/k}{P_{\text{noise}}/k} = \frac{P_{\text{off}}}{P_{\text{noise}}} = \frac{P_{\text{off}}}{P_{\text{on}} - P_{\text{off}}}, \quad (3.8)$$



(a)



(b)

Figure 3.3: *Visibilities before (a) and after (b) inverting the gain corrections for a typical observation and a typical subband, baseline, and polarization. Tapering was already applied to both datasets (see Sect. 3.1). The bright peaks are caused by RFI. In (a) dark horizontal bands are visible due to stronger gain correction. The color palette of appendix A.2 is used for displaying complex values. The domination of red in the figures suggests a big real contribution in the complex visibilities.*

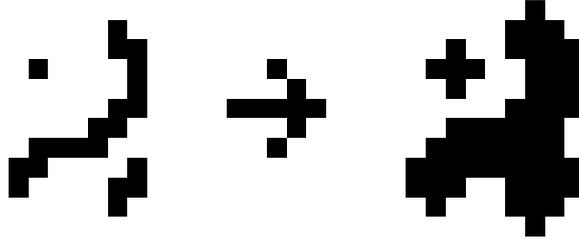


Figure 3.4: *The effect of the dilation algorithm. The original image (left hand side) becomes fattened (right hand side).*

and therefore:

$$T_{\text{sys}} = T_{\text{noise}} \frac{P_{\text{off}}}{P_{\text{on}} - P_{\text{off}}}. \quad (3.9)$$

Here P_{on} and P_{off} are the power measurements with and without added noise, respectively, and k is the Boltzmann constant. Equations (3.8) and (3.9) are only valid if the internal gain is constant during the two power measurements.

The two power measurements P_{on} and P_{off} are both affected by RFI. The RFI in successive P_{on} and P_{off} measurements is mostly uncorrelated (see Fig. 3.2). This introduces errors in the T_{sys} measurements. Also, P_{off} measurements have a higher signal to noise ratio than P_{on} measurements. Measuring P_{on} and P_{off} takes one and five to six seconds, respectively. P_{off} is therefore more representative of the average conditions during an observation than P_{on} or T_{sys} .

To get an absolute flux scale both the T_{sys} and the absolute gain of each antenna, polarization, and subband should be known. Unfortunately we do not have a model for the absolute gains of the WSRT/LFFE. Transforming the unitless P_{off} values to T_{sys} values is therefore unnecessary and inversion of the AGC gain correction was done using P_{off} instead of T_{sys} .

The correlator uses voltages as input and not powers. Therefore inverting the AGC gain corrections was done by taking the square-roots of the corresponding signal powers like:

$$\mathbf{V}'_{pq} = \sqrt{\begin{pmatrix} a_x & a_y \end{pmatrix}_p} \begin{pmatrix} xx & xy \\ yx & yy \end{pmatrix}_{pq} \sqrt{\begin{pmatrix} a_x \\ a_y \end{pmatrix}_q}, \quad (3.10)$$

where \mathbf{V}'_{pq} is the observed visibility of baseline pq after inversion of the gain corrections, a_x is the P_{off} in the x polarization of antenna p and the central matrix is the observed visibility on baseline pq before inversion of the gain corrections.

The effects of typical gain corrections is shown in Fig. 3.3(a). Occurring RFI increases T_{sys} and therefore the AGC suppresses the signal for the whole subband. When RFI appears in just a few channels the signal in all channels decreases. This causes the dark horizontal bands in Fig. 3.3(a). Fig. 3.3(b) shows the visibilities after compensating for

the gain corrections. The dark bands have disappeared.

Flagging data should be done after inverting the gain corrections. If flagging was done before inverting gain corrections, visibilities in the dark bands could become flagged due to their lower amplitude.

3.3 Data quality

In the previous two sections we have done two necessary pre-processing steps. We are now ready to verify the quality of our observations and remove unreliable data.

In this section we will see that a large fraction of our data has been corrupted. There are two causes for corruption: problems with (parts of) the WSRT and interference. Corruption can take place on many levels such as visibilities, frequency channels, antennas, or even entire observation sessions. In Sect. 3.4 we use an algorithm to identify bad data on a deep level: visibilities and baselines. However, processing on a deep level is inefficient if data has been corrupted at a much higher level. This section deals with identifying bad data at these higher levels and excluding them from further analysis.

We operate at low radio frequencies in heavily used, unprotected frequency bands. Therefore we have a lot of RFI corrupting our data. Some of the RFI is Earth-bound and some comes from the Sun. There is RFI in all frequency bands and during all observing sessions. Figure 3.5 shows visibilities for one observation on a short baseline during typical observing conditions. The bright peaks are RFI.

In order to not complicate the data reduction more than strictly necessary, the subbands that are least affected by RFI should be selected for further analysis. From visual inspections (see for example Fig. 3.5), we have chosen the fourth subband for further analysis. Therefore the data analysis in the remainder of this report only deals with the fourth subband of our observations. The central frequency of the fourth subband is 142.0 MHz (see Tab. 2.1). A detail of Fig. 3.5 for the fourth subband is shown in Fig. 3.6.

Not all observations in the fourth subband are usable. A way to identify them is by looking at deviating visibility amplitudes. Figure 3.9(a) shows visibility amplitudes (the amplitude median of channels 10-100 of the fourth subband, polarization xx) for five baselines and all the time slots of field 3C147. The time slot labels separate the eight sessions (3C147 was observed during all eight sessions). The figure shows:

- Different amplitudes occur during the entire first session (so not only during the 90 time slots of the single 3C147 observation). The strange low amplitudes are not caused by RFI, but probably by different settings of the WSRT backend. The first session is therefore unusable and excluded from further analysis.
- Different amplitudes occur for baseline 0/D (blue curve) during the entire second session. This also happens for baseline 3/7 (magenta curve) during part of the fourth session. Inspection of other baselines shows that there is something wrong

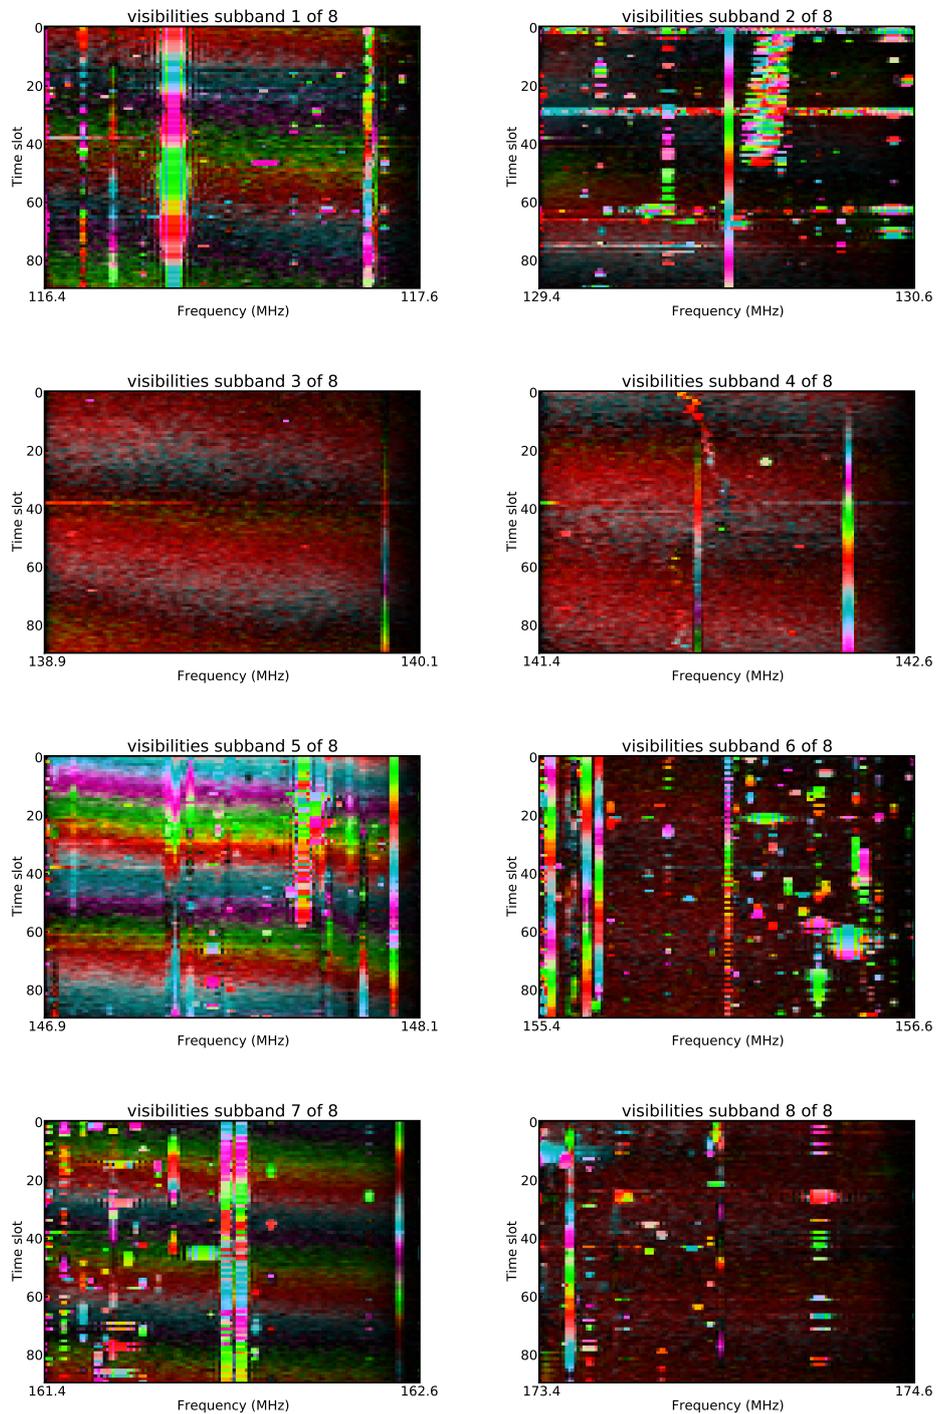


Figure 3.5: *Typical visibilities for all eight subbands for baseline 1/2 and polarization xx during one observation of 3C147, after tapering and inversion of the gain correction. There are frequency gaps between the eight subbands. The fourth subband in this figure is shown in detail in Fig. 3.6. The color palette of appendix A.2 is used for displaying complex values.*

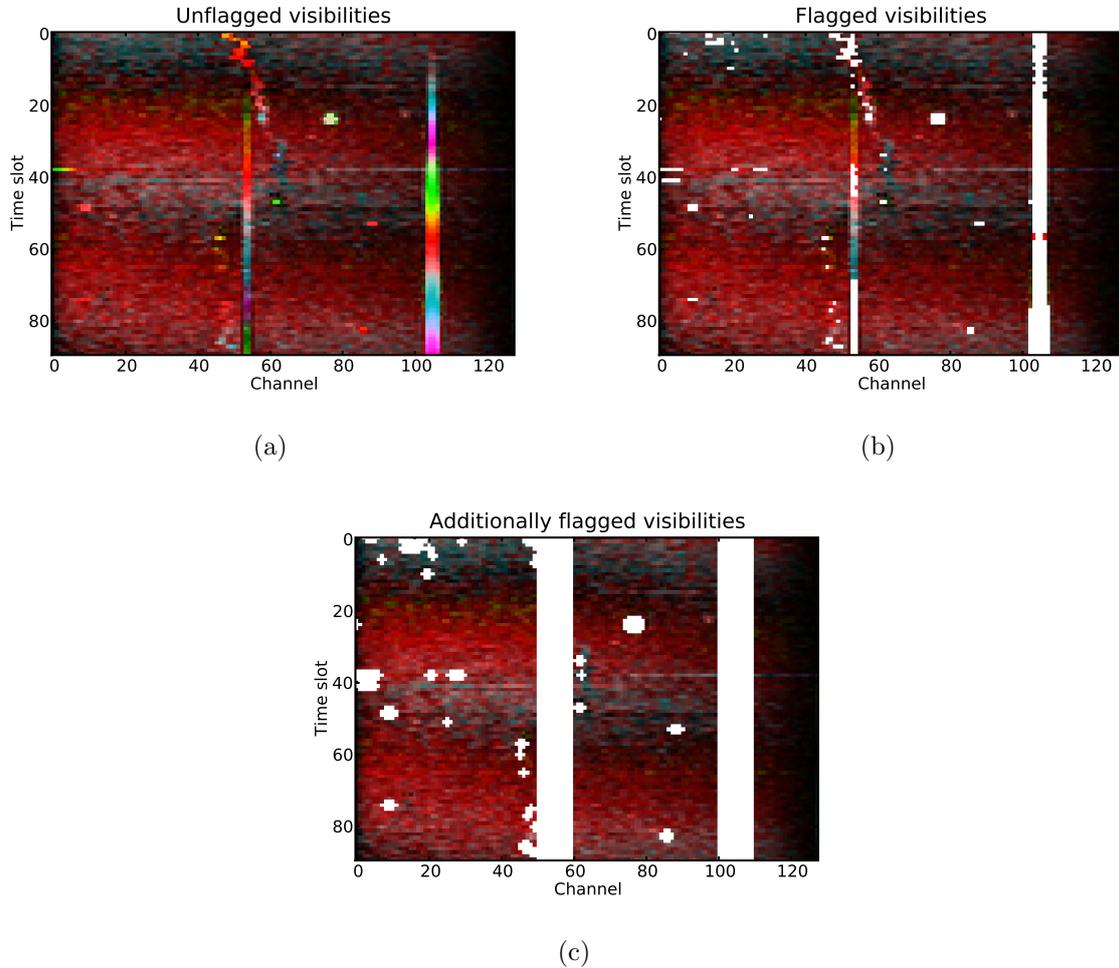


Figure 3.6: *Typical visibilities for one observation of 3C147 and for the fourth subband, baseline 1/2, and polarization xx , after tapering and inversion of the AGC. This figure is a detail of Fig. 3.5. In (a) the visibilities are shown before flagging any RFI. In (b) the visibilities are shown after flagging using the WSRT flagger algorithm in Sect. 3.4. Flagged visibilities are shown in white. In (c) the visibilities are shown after final flagging in Sect. 3.5: a dilation algorithm (explained in Fig. 3.4) was applied and channels dominated by RFI were manually flagged (channels 50-60 and 100-110 for this observation). The color palette of appendix A.2 is used for displaying complex values.*

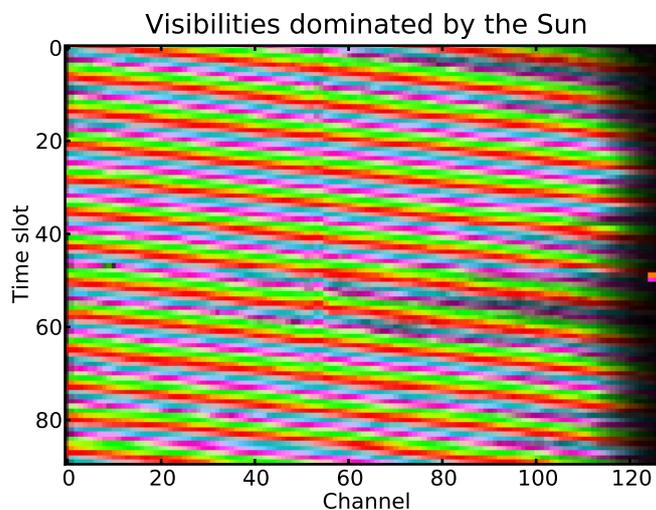


Figure 3.7: *Visibilities of a typical observation, subband, baseline, and polarization during the fifth session. The strong fringes are caused by temporary radiation from the Sun. The color palette of appendix A.2 is used for displaying complex values.*

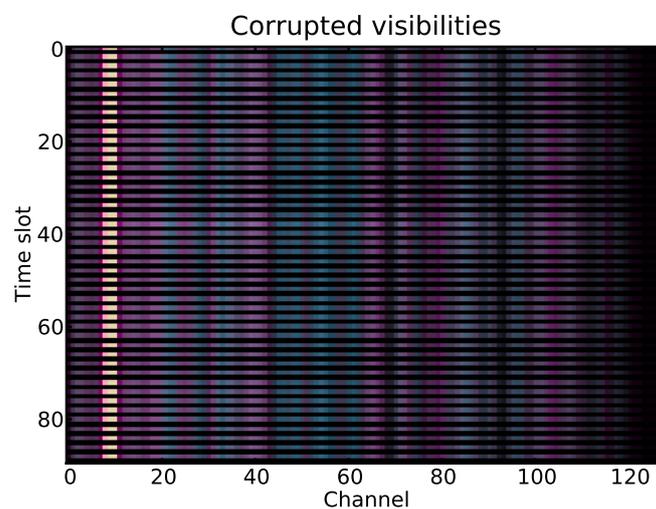
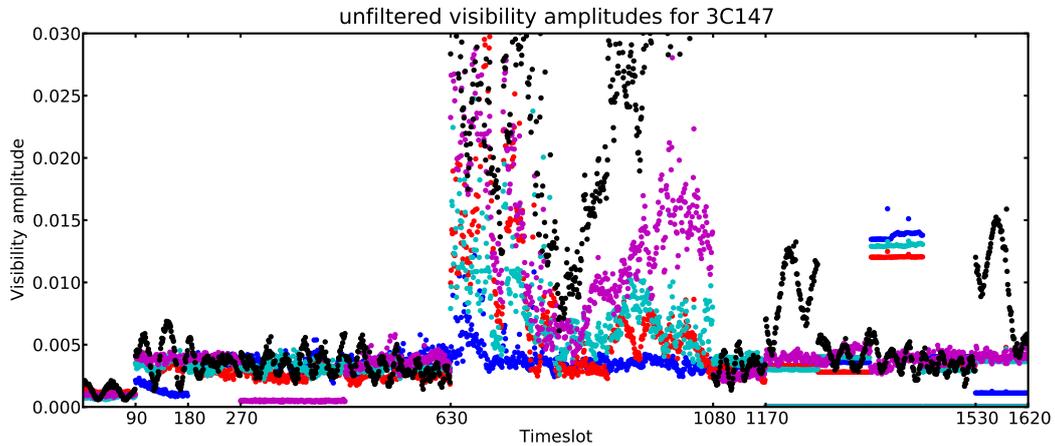
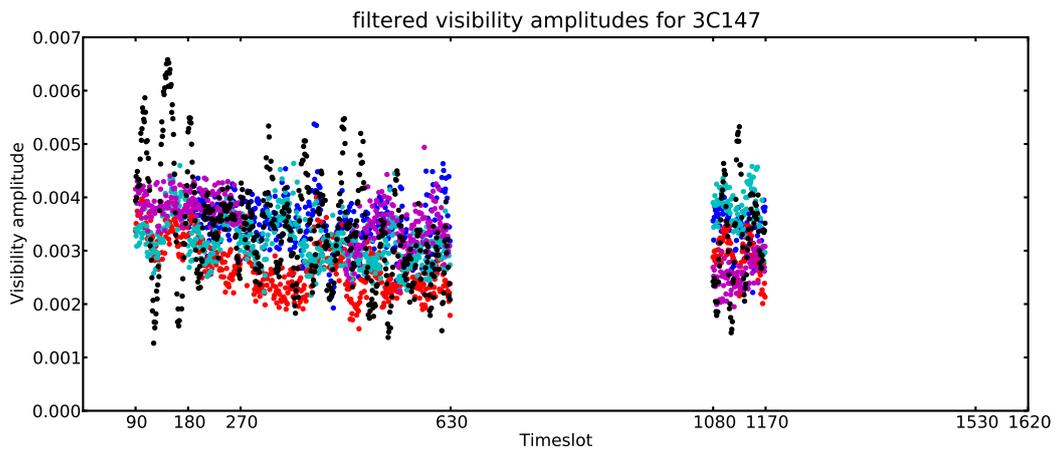


Figure 3.8: *Visibilities of a typical observation of the seventh and eighth session for the fourth subband and a typical baseline and polarization. The strange visibilities are probably caused by malfunctioning of the correlator. The color palette of appendix A.2 is used for displaying complex values.*



(a)



(b)

Figure 3.9: *Visibility amplitudes of all the observations of 3C147 for the fourth subband, for five baselines, and the xx polarization. For each time slot the median amplitude of channels 10-100 was taken. (a) shows the amplitudes for all the visibilities and (b) shows them after dropping sessions and flagging visibilities, antennas, and channels as described in Sect. 3.3, Sect. 3.4, and Sect. 3.5. Figure 3.6 corresponds to time slots 90-180. The time slot labels separate the eight sessions. The colors blue, red, cyan, magenta, and black correspond to baselines 0/D, 1/B, 2/9, 3/7, and 4/5.*

with one antenna (antenna D during the first session and antenna 3 during the first 18 observations of the third session). Therefore flagging antennas on deviating behavior is necessary. This is described in Sect. 3.5.

- Different amplitudes occur during the entire fifth session. Figure 3.7 shows typical visibilities during this session. The strong fringes are caused by temporary radiation from the Sun: during the observing session at 13/08/2006 8:45 UT a bright flare appears in X-rays (Freeland, 2009), as detected by the Goes 12 Solar X-Ray Imager (Hill et al., 2005). A few minutes later the strong fringes appear in our radio observations. The longest baseline 0/D (the blue curve in Fig. 3.9(a)) is not corrupted by the solar radiation. The radiation source at the Sun has big enough celestial dimensions to be invisible on the longest baseline. With only a few usable baselines self calibration will not succeed. The fifth session is therefore unusable and excluded from further analysis.
- Different amplitudes occur for multiple baselines during the entire seventh and eighth session. Fig. 3.8 shows typical visibilities from this session. Inspection showed that the deviating visibilities occur for all baselines where one antenna is in the 0-7 range and the other is in the 8-D range. Within these ranges the visibilities are not deviating (such as baseline 4/5, the black curve). The strange visibilities are probably caused by malfunctioning of the correlator. With only a few usable baselines self calibration will not succeed. The seventh and eighth sessions are therefore unusable and excluded from further analysis.

The first, fifth, seventh, and eighth session are excluded from further analysis and four sessions remain. From now on, these four sessions are called the first, second, thirth, and fourth session.

Of the four remaining sessions, three of the in total 96 observations do not span 90 time slots. Two of them are of Cyg A and one is of Cas A. None of them is unique for its field during its session. Therefore they do not add important hour angles for improving the qualities of their field images. Because it is very convenient that all observations have the same number of time slots (especially during self calibration), the three observations were excluded from further analysis.

3.4 WSRT flagger

After identifying bad data on a high level in the previous section we are now ready to identify bad data on a deeper level. Figure 3.6(a) shows typical visibilities and all the RFI visible here should be flagged. This can be done with the WSRT flagger developed by Renting and Brentjens (2006). The algorithm was originally developed for higher frequencies, but can also be used at lower frequencies. The algorithm flags visibilities with directions in the complex plane that deviate too much from a moving median window in

time. The algorithm can also flag baselines of a time slot if the root mean square (RMS) of its unflagged visibility amplitudes deviate. The RMS should not deviate more than a factor two of the noise estimate, which is $1/\sqrt{t_{\text{integration}} \cdot \text{bandwidth}_{\text{channel}}}$. The RMS should also not deviate more than a factor ten of other baselines RMSs. This way, for example, malfunctioning antennas get flagged.

Two settings of the algorithm should be adapted to the used datasets: the window size and threshold. The best results were accomplished using a large window size of 40 time slots. Since the visibilities of one channels look rather repetitive over time (see for example Fig. 3.6), the median of the windowed visibilities becomes a more or less fixed point in complex space if the window size is large. Therefore the flagger basically becomes a limiter: all visibilities with distances from the fixed point that are larger than the threshold get flagged.

The selected threshold depends on the observed field. The manually selected thresholds for the fields were: 10 for Cyg A, 8 for Cas A, and 1.5 for all the other fields.

Figure 3.6(b) shows typical flagging results. Not all RFI was flagged. Changing the settings did not improve this: RFI remained unflagged while 'good' data got more and more flagged. Improving the flagging is investigated in the next section.

3.5 Additional flagging

After flagging visibilities in the previous section we can investigate the results in two ways. We can look at our observations at different levels and reject parts that have too many flagged visibilities. We can also look at the visibilities in detail and judge how flagging was insufficient and should be augmented.

The fractions of all the visibilities that are flagged can be investigated on different levels:

- Per observation and antenna, as shown in Fig. 3.10. The observation labels separate the four sessions. For some sessions the fractions flagged are high for one antenna (antenna D during the first session and antenna 3 during the first 18 observations of the third session). Here, the antennas were not working as expected. This was already mentioned in Sect. 3.3. The malfunctioning antennas were fully flagged during the observations with deviating behavior.
- Per observation and field, as shown in Fig. 3.11. The fractions flagged are below 30% for most observations and change with session. The fractions are very low during the second session and the first half of the last session. The fractions do not depend on the observed field. Therefore the observed RFI shows no direction dependency (since the fields have very different celestial positions). Figure 3.11 shows two observations with high fractions flagged: one for Cas A and one for Cyg A. These observations are unusable because probably all visibilities are more or less affected by RFI and the

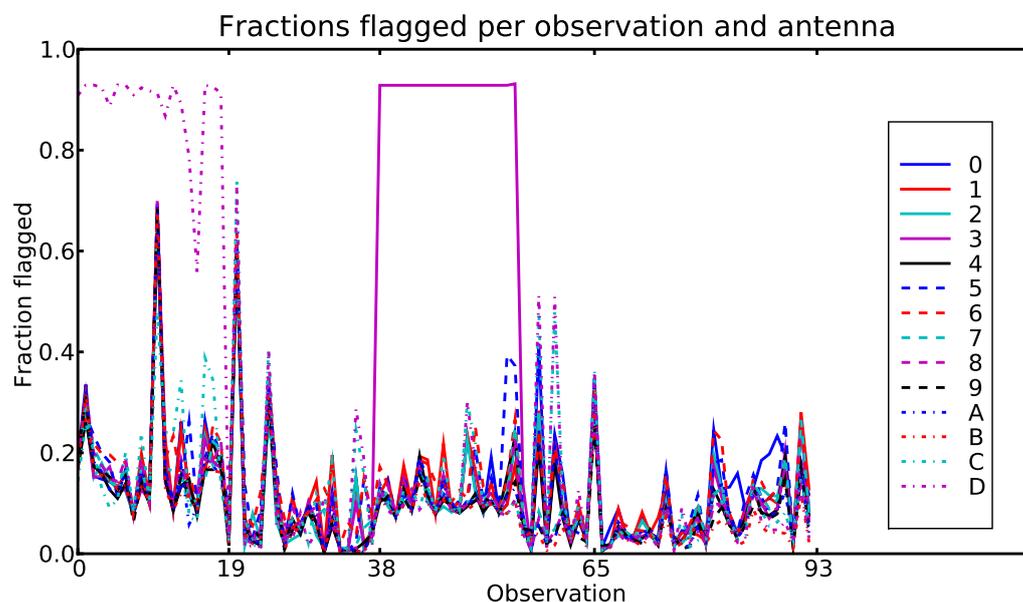


Figure 3.10: *The fractions of the visibilities that are flagged per observation and antenna. Only visibilities of the fourth subband were used. The observation labels separate the four sessions. Antenna D was not working properly during the first session and antenna 3 not during the first 18 observations of the third session.*

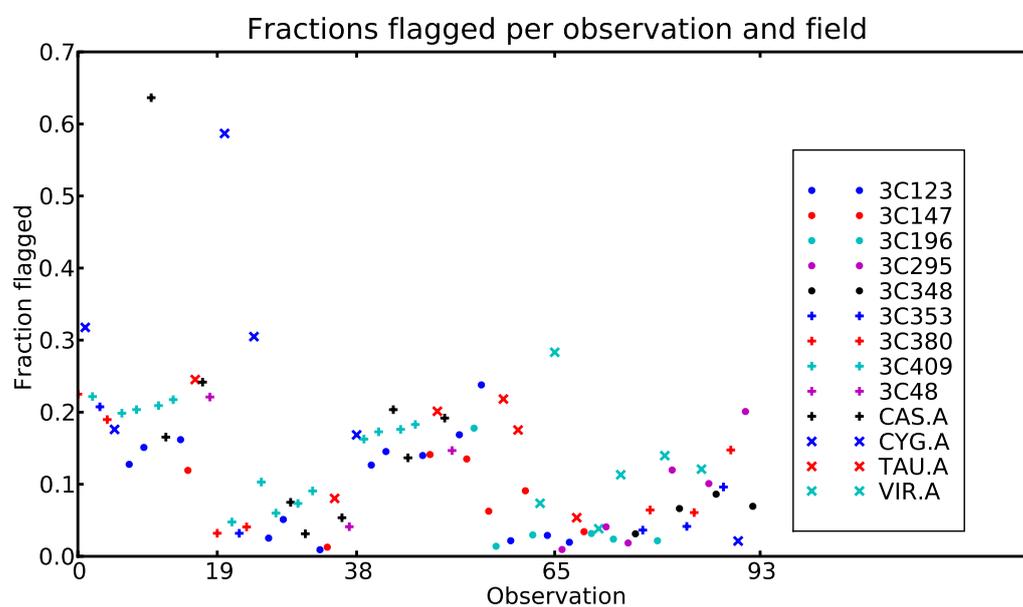


Figure 3.11: *The fractions of the visibilities that are flagged per observation and field. Only visibilities of the fourth subband were used. The observation labels separate the four sessions.*

RFI frustrates self calibration. Therefore the two observations were excluded from further analysis.

For judging how flagging was insufficient, Fig. 3.6(b) shows typical flagging results. The flagged visibilities are shown in white. The figure shows typical flagging shortcomings:

- Some flagged regions still show RFI surrounding it. The edges of RFI peaks are ignored by the flagger algorithm. Flagging the edges can be done by applying a dilation algorithm (explained in Fig. 3.4) on the flagged visibilities. The result is shown in Fig. 3.6(c).
- Some channels show RFI during the whole observation that was not fully flagged. The flagging failure can be due to the used moving median window in time: all the visibilities delivered by the window are corrupted. Many of our observations show channels dominated by RFI. Flagging those channels was done manually. For one observation the result is shown in Fig. 3.6(c), where channels 50-60 and 100-110 were manually flagged. Inspecting all the visibilities is unpractical since there are 91 observations, 91 baselines, and 2 polarizations. However, RFI typically appears simultaneously for all baselines and polarizations. It was enough to inspect only one baseline and polarization for each observation and apply the flagging of found bad channels to all baselines and polarizations of the observation.
- Some time slots show hints of deviating visibilities appearing over a big range of channels. For example at time slot 48 in Fig. 3.6(b). These hints are very subtle and do not appear at other polarizations or baselines. We decided to leave these time slots intact. Therefore we did not classify this as a flagging shortcoming.

We have now finished verifying the quality of our observations and decided what data is acceptable for further analysis. The deviating visibility amplitudes of Fig. 3.9(b) have all disappeared for our new selection of data in Fig. 3.9(b).

3.6 Scale visibility amplitudes per field

During self calibration the observations are compared with a model. Part of this model is the local sky model (LSM), a model of the Stokes parameters and celestial positions of the sources in the field-of-view. The Stokes parameters in the LSM will be fit to the unitless fluxes in our observations. However, there is no reason to believe that the compared quantities have the same scaling. To simplify the fitting we can adapt our unitless observations or the LSM in advance. Here we decided to scale the observations of each field. This has one important consequence: the relative fluxes measured from the self calibrated observations in Sect. 5.1 should be compensated for the scaling factors applied here.

When scaling our observations we have to know how to retrieve the unitless fluxes from our data. At this stage we know only one thing about our fluxes: we expect that observations are dominated by the main source we pointed the antennas at and/or the Galactic

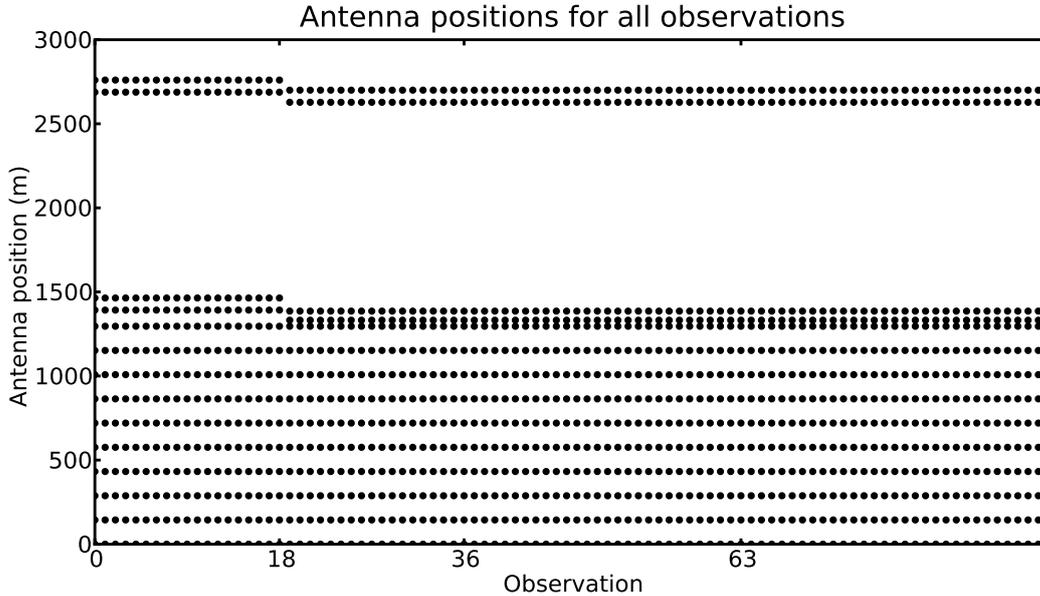


Figure 3.12: *Antenna distances to antenna 0 for all the observations of the four selected sessions. Antennas A-D have been moved between the first and second session.*

sky background. And this simplifies getting a first glance at the fluxes a lot. If the main source is constantly in the phase center it is the only source having a 100% real contribution to the complex visibilities for all hour angles. The average real part of all unflagged visibilities would then give us an estimate of the flux of the main source (more on this in Sect. 5.1). However, due to ionospheric phase shifts the main source will move out of the phase center and its visibility contribution will not be 100% real. When the main source dominates observations the flux of the main source can then be estimated by averaging the amplitudes of all unflagged visibilities instead of the real parts. The amplitudes are less affected by ionospheric phase shifts than the real parts.

Using the flux estimate of the main source, the observations of each field can be scaled so the flux estimate of the main source becomes 1. The applied scaling values for all the fields are given in Tab. 3.1. The scaling values highly depend on the field observed. Especially the A-team observations needed little scaling.

3.7 Renaming moved antennas

In the next chapter self calibration is described for all fields using calibration package MeqTrees. For each field all the observations are combined to one data set that will be used by the software. However, Fig. 3.12 shows that antennas A-D have been moved between the first and second session. Therefore within a single data set antennas A-D can

have two antenna positions. This is not expected by MeqTrees. As a solution for this, the antennas A-D were renamed to E-H for the observations of all but the first session. However, in the remainder of this report, we continue to call them A-D.

Table 3.1: *Scaling factors applied to all observations of each field. Scaling was done to prepare for self calibration using a simple initial model.*

Name	Scaling factor
3C48	336.53
3C123	89.76
3C147	347.08
3C196	292.40
3C295	272.99
3C348	72.81
3C353	118.33
3C380	291.71
3C409	215.65
Cas A	3.78
Cyg A	2.32
Tau A	23.09
Vir A	38.59

Chapter 4

Self calibration

Observations of astronomical objects by any instrument at the Earth's surface are influenced by instrumental effects and the ionosphere. An effective way of dealing with these effects is by doing a self calibration (Pearson and Readhead, 1984). In this procedure, corrections for instrumental and atmospheric effects are derived based on a model of the sky, a model of the instrumental and atmospheric effects, and the observed Stokes parameters (Hamaker et al., 1996; Sault et al., 1996; Hamaker, 2000).

In this chapter the complete self calibration process is described and applied to our observations. It is an important step done before analyzing the data in chapter 5. The process is first described by formulating all the instrumental effects and their unknowns using the measurement equation.

Formulating the measurement equation is done in two phases. It is first formulated for the WSRT in general. Afterwards, we discuss how we can simplify it for our observations and goals.

Before we can follow all the steps in self calibrating we have to obtain an initial local sky model (LSM). Only with this model we can make a first guess of what an observation looks like. We have chosen to take LSMs of point sources. Therefore six sources (3C348, 3C353, Cas A, Cyg A, Tau A, and Vir A) were dropped from our project because they are clearly extended.

Setting up the self calibration was done using the software package MeqTrees (Noordam and Smirnov, 2009). Our measurement equation is implemented using a Python script and run by MeqTrees. With each run, different parameters were set to do different calibrations. For each field six steps of calibration were done: bandpass, phase, flux, a repetition of phase and flux, and finally gain calibration. Doing flux calibration (calibrating source Stokes parameters and celestial positions) was complicated: due to memory limits MeqTrees was not able to do this calibration on a complete data set of a field. We attempted cutting datasets into pieces and combine the results from each part. The gain calibration gave unsatisfying results, giving a hint that our LSMs were not sufficient.

When all calibrations were done the quality of the calibration was verified by inspecting the residuals: the differences between observed visibilities and model visibilities. For some fields there was a significant difference between the main source in the observation and

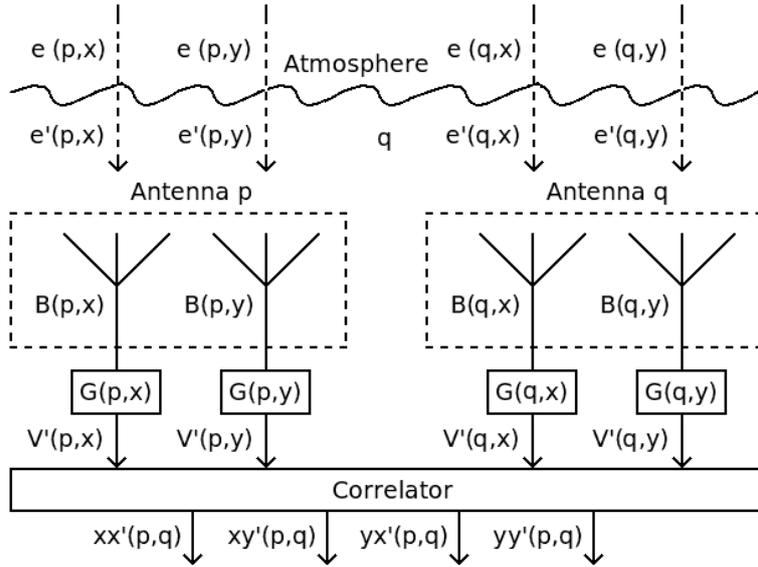


Figure 4.1: *Block diagram of one baseline of a linear polarization interferometer like the WSRT, which is a linear version of the block diagram presented in Hamaker et al. (1996). $e(p,x)$ = real electric field for antenna p and polarization x , $e'(p,x)$ = observed electric field, $B(p,x)$ = antenna beam, $G(p,x)$ = antenna gain, $V'(p,x)$ = observed voltage, $xx'(p,q)$ = observed visibility for the x polarizations of antennas p and q .*

the model. It was visible in the residual. Therefore the self calibration was imperfect. Unfortunately due to time constraints we were not able to search for solutions.

4.1 Formulating the measurement equation for the WSRT

For self calibration we need a model of the observations. This model should describe how the electric field from the sky is recorded by the instrument. In the case of the WSRT, the instrument is a linear polarization interferometer. The block diagram of one baseline of such an instrument is shown in Fig. 4.1, which is a linear version of the block diagram presented in Hamaker et al. (1996). The electric field is transformed several times before the final recordings are done. The model for the observations should contain all the transformations. In this section, the model for the observations is formulated, as described in detail in a series of five papers: Hamaker et al. (1996); Sault et al. (1996); Hamaker and Bregman (1996); Hamaker (2000, 2006).

The electric field from the sky is transformed by the ionosphere and the instrument. If the transformation is linear, it can be described as a 2×2 complex matrix, known as a Jones

Matrix \mathbf{J} , acting on the arriving electric field:

$$\mathbf{e}' = \mathbf{J}\mathbf{e}, \quad (4.1)$$

where \mathbf{e} is the real electric field and \mathbf{e}' is the observed electric field. Both \mathbf{e} and \mathbf{e}' are 2D vectors with components for two orthogonal polarizations.

For an interferometer the electric field enters multiple antennas and polarizations, each introducing its own instrumental and atmospheric effects. The observed visibility \mathbf{V}'_{pq} for antennas p and q for a source on the phase center is given by:

$$\mathbf{V}'_{pq} = \langle \mathbf{e}'_p \mathbf{e}'_q \rangle = \langle (\mathbf{J}_p \mathbf{e})(\mathbf{J}_q \mathbf{e})^\dagger \rangle = \langle \mathbf{J}_p (\mathbf{e} \mathbf{e}^\dagger) \mathbf{J}_q^\dagger \rangle = \mathbf{J}_p \langle \mathbf{e} \mathbf{e}^\dagger \rangle \mathbf{J}_q^\dagger = \mathbf{J}_p \mathbf{B} \mathbf{J}_q^\dagger, \quad (4.2)$$

where the \dagger represents the conjugate transpose matrix operator and \mathbf{J}_p is the Jones matrix for antenna p which is assumed constant during the correlation interval. \mathbf{B} is the coherency matrix of the source:

$$\mathbf{B} = \langle \mathbf{e} \mathbf{e}^\dagger \rangle = \frac{1}{2} \begin{pmatrix} I + Q & U - iV \\ U + iV & I - Q \end{pmatrix}, \quad (4.3)$$

where I, Q, U , and V are the Stokes parameters defining both intensity and polarization of the source.

With sources away from the phase center we need to define coordinate systems for both the antennas and the sources on the sky. The coordinates of an antenna (u, v, w) are measured in a right-handed coordinate system, where u and v are measured in a plane normal to the direction of the phase center (Thompson et al., 2001). The component v is measured towards the north as defined by the plane through the origin, the phase center, and the pole. The component u is measured towards the east. The component w is measured in the direction of the phase center. All three components can be measured in length units and wavelengths. Baselines are derived by subtracting antenna coordinates:

$$(u_{pq}, v_{pq}, w_{pq}) = (u_p - u_q, v_p - v_q, w_p - w_q). \quad (4.4)$$

The coordinates of a source on the sky (l, m) is measured relative to the phase center, where l and m are defined as the cosines of the angles (the *direction cosines*) between the direction (l, m) and the u and v axes, respectively.

With N sources having coherencies B_s and celestial positions (l_s, m_s) , the observed visibility of baseline (u_{pq}, v_{pq}, w_{pq}) is given by:

$$\mathbf{V}'_{pq} = \sum_{s=1}^N \mathbf{J}_{p,s} \mathbf{B}_s \mathbf{J}_{q,s}^\dagger e^{2\pi i(u_{pq}l_s + v_{pq}m_s + w_{pq}(n_s - 1))\frac{z}{c}}, \quad (4.5)$$

where $n_s = \sqrt{1 - l_s^2 - m_s^2}$ are the direction cosines measured relative to the w axis. Equation 4.5 is an expansion of Eq. (4.2) with added phase terms for all the sources. The Jones matrices $\mathbf{J}_{p,s}$ are source dependent because the instrumental and atmospheric effects can

be direction dependent.

The position of the phase center depends on the applied fringe stopping: an instrument like the WSRT in general tracks a celestial position in pointing and delay while the Earth rotates.

Using Eq. (4.4), Eq. (4.5) can be rearranged in order to separate antenna dependent components:

$$\begin{aligned} \mathbf{V}'_{pq} &= \sum_{s=1}^N e^{2\pi i(u_p l_s + v_p m_s + w_p(n_s-1))\frac{z}{c}} \mathbf{J}_{p,s} \mathbf{B}_s \mathbf{J}_{q,s}^\dagger e^{-2\pi i(u_q l_s + v_q m_s + w_q(n_s-1))\frac{z}{c}} \\ &= \sum_{s=1}^N \mathbf{K}_{p,s} \mathbf{J}_{p,s} \mathbf{B}_s \mathbf{J}_{q,s}^\dagger \mathbf{K}_{q,s}^\dagger, \end{aligned} \quad (4.6)$$

where $\mathbf{K}_{p,s}$ is a phase shift term, a scalar Jones matrix:

$$\mathbf{K}_{p,s} = \begin{pmatrix} e^{2\pi i(u_p l_s + v_p m_s + w_p(n_s-1))\frac{z}{c}} & 0 \\ 0 & e^{2\pi i(u_p l_s + v_p m_s + w_p(n_s-1))\frac{z}{c}} \end{pmatrix}. \quad (4.7)$$

Not all instrumental and atmospheric effects are direction dependent (like the gain of an antenna). Therefore $\mathbf{J}_{p,s}$ can be split into components:

$$\mathbf{J}_{p,s} = \mathbf{G}_p \mathbf{E}_{p,s}. \quad (4.8)$$

These terms can be split further in time and/or frequency dependent terms such as antenna gain and antenna passband. The visibility equation becomes:

$$\mathbf{V}'_{pq} = \mathbf{G}_p \left[\sum_{s=1}^N \mathbf{E}_{p,s} \mathbf{K}_{p,s} \mathbf{B}_s \mathbf{K}_{q,s}^\dagger \mathbf{E}_{q,s}^\dagger \right] \mathbf{G}_q^\dagger. \quad (4.9)$$

Because $\mathbf{K}_{p,s}$ is a scalar Jones matrix it commutes with both $\mathbf{E}_{p,s}$ and \mathbf{G}_p .

The equation just presented is an example of a measurement equation. It gives the relation between the measurements by the interferometer, the electric fields from the celestial objects, and all the transformations in between.

4.2 Measurement equation for our observations

During self calibration one attempts to solve for subsets of the parameters of the measurement equation until the values of all parameters are found. This may require several iterations, solving for different subsets of parameters in each step. This section summarizes all the measurement equation components in our observations and their corresponding

transformations and parameters.

The LSM is a collection of Stokes parameters and celestial positions for all the sources in the field-of-view. In the measurement equation 4.9 they were called \mathbf{B}_s and (l_s, m_s) , respectively. The Stokes parameters and celestial positions can be the parameters of the LSM. However, we will assume that all sources are unpolarized and only need Stokes I . Therefore, the LSM parameters are fluxes and celestial positions. We consider them constant with time (no variable or moving sources). Obtaining the initial LSMs for all fields is described in Sect. 4.3.

As discussed in Sect. 4.1, instrumental effects can be described by 2×2 complex matrices. These matrices contain four elements with indices xx , xy , yx , and yy , which contain the parameters of an instrumental effect.

However, for our observations only calibration for instrumental xx and yy components is required. The sources are assumed unpolarized and only xx and yy visibilities were stored for our observations. The xy and yx visibilities were not determined by the WSRT correlator. The next paragraph shows why only calibration for instrumental xx and yy is required here.

According to Eq. (4.3) the coherency of an unpolarized source with total flux F is:

$$\mathbf{B} = \begin{pmatrix} b_{xx} & 0 \\ 0 & b_{yy} \end{pmatrix}, \quad (4.10)$$

where $b_{xx} = b_{yy} = \frac{1}{2}F$. However, we will not rename b_{xx} and b_{yy} in the following steps. If the source is in the phase center of a linear polarization interferometer, then according to Eq. (4.2) the observed visibility for antennas p and q is:

$$\mathbf{V}'_{pq} = \mathbf{J}_p \mathbf{B} \mathbf{J}_q^\dagger, \quad (4.11)$$

or:

$$\begin{pmatrix} xx' & xy' \\ yx' & yy' \end{pmatrix}_{pq} = \begin{pmatrix} p_{xx} & p_{xy} \\ p_{yx} & p_{yy} \end{pmatrix} \begin{pmatrix} b_{xx} & 0 \\ 0 & b_{yy} \end{pmatrix} \begin{pmatrix} q_{xx}^* & q_{yx}^* \\ q_{xy}^* & q_{yy}^* \end{pmatrix}, \quad (4.12)$$

or:

$$\begin{pmatrix} xx' & xy' \\ yx' & yy' \end{pmatrix}_{pq} = \begin{pmatrix} p_{xx} & p_{xy} \\ p_{yx} & p_{yy} \end{pmatrix} \begin{pmatrix} b_{xx}q_{xx}^* & b_{xx}q_{yx}^* \\ b_{yy}q_{xy}^* & b_{yy}q_{yy}^* \end{pmatrix}, \quad (4.13)$$

or:

$$\begin{pmatrix} xx' & xy' \\ yx' & yy' \end{pmatrix}_{pq} = \begin{pmatrix} p_{xx}b_{xx}q_{xx}^* + p_{xy}b_{yy}q_{xy}^* & p_{xx}b_{xx}q_{yx}^* + p_{xy}b_{yy}q_{yy}^* \\ p_{yx}b_{xx}q_{xx}^* + p_{yy}b_{yy}q_{xy}^* & p_{yx}b_{xx}q_{yx}^* + p_{yy}b_{yy}q_{yy}^* \end{pmatrix}, \quad (4.14)$$

where visibility components xy' and yx' were not recorded. In the matrices p and q are the components for the instrumental effects of antennas p and q , respectively. Equation 4.14 shows a leakage from b_{yy} to xx' and from b_{xx} to yy' and these depend on the xy and yx terms of the instrumental effects. However, these off-diagonal terms have an expected power amplitude of a few percent. Therefore the leakage from one recorded polarization to

the other is negligible. As a result, the xy and yx terms of the instrumental effects do not have to be solved for during self calibration. If solving for polarizations of instrumental effects is mentioned in the next paragraphs, solving is only done for individual polarizations (xx and yy) and not for cross leakage terms (xy and yx).

Each antenna and polarization of the WSRT backend has its own passband filter. The effects of a passband filter can be seen in Fig. 3.6(a): for the highest channels the intensity is lower, caused by the lower transmission of the bandpass at these channels. Besides an amplitude effect, the bandpass filter also introduces a phase shift.

Amplitude and phase are frequency dependent and can be time dependent as well. However, they do not depend on the field being observed: it is a strictly internal instrumental effect. Therefore if the filter characteristics can be considered constant over a time interval, one set of solutions can be used for all observations within the time interval.

Bandpass calibration is the only frequency dependent effect solved for. Bandpass calibration is discussed in detail in Sect. 4.5.

Each antenna has its own beam: its sensitivity depends on the angle from the antenna axis. Due to this effect, off-axis sources appear weaker than expected. However, in this project we are not interested in the real fluxes of off-axis sources. We just want to subtract them from our observations so only the central source remains (more on this in Sect. 5.1). If we presume that the beam is constant with time (the WSRT antennas are equatorially mounted) and the same for each antenna and polarization, we can ignore the beam and leave them out of our measurement equation. However, we will find this assumption is wrong: the antenna beam strongly depends on polarization (see Sect. 4.10).

Radio observations are affected by the ionosphere. Especially in our frequency regime the ionosphere is very important. It shifts the phases with high speed variations but leaves the amplitudes more or less unchanged. For this reason, the integration time of each observation was only 10 seconds allowing us to solve for ionospheric phase shifts. Compensating for ionospheric phase shifts is very important, because phase shifts move and distort the sources in the image plane.

The parameter phase depends on the antenna and polarization. The ionosphere is different for different positions on the Earth's surface. Within the size of the WSRT, the phases are only expected to have a slope with increasing baseline. The phase varies with time, but does not vary within our frequency subband (141.45 MHz to 142.55 MHz). Phase calibration is discussed in detail in Sect. 4.6.

Besides a bandpass filter, each antenna and polarization has a time dependant amplifier gain. We assume this is strictly an amplitude effect. If it also shifts the phase, this effect cannot be separated from the ionospheric phase shift. If the shape of the bandpass amplitude is time independent, but its overall amplitude is not, this is the result of a time dependent, frequency independent antenna gain. Gain calibration is discussed in detail in Sect. 4.9.

Taking all the effects into account, our measurement equation is:

$$\mathbf{V}'_{pq}(\nu, t) = \mathbf{P}_p(\nu) \mathbf{G}_p(t) \mathbf{A}_p(t) \left[\sum_{s=1}^N \mathbf{K}_{p,s}(\nu, t) \mathbf{B}_s \mathbf{K}_{q,s}^\dagger(\nu, t) \right] \mathbf{A}_q^\dagger(t) \mathbf{G}_q^\dagger(t) \mathbf{P}_q^\dagger(\nu), \quad (4.15)$$

where \mathbf{P}_p is the antenna bandpass, \mathbf{G}_p is the antenna gain, \mathbf{A}_p is the ionospheric phase shift for an antenna, $\mathbf{K}_{p,s}$ is the interferometric phase shift for a source and antenna, and \mathbf{B}_s is the coherency matrix of a source. All the identities are diagonal matrices containing only xx and yy components. And since these diagonal matrices commute, the order of the components in Eq. (4.15) can be exchanged.

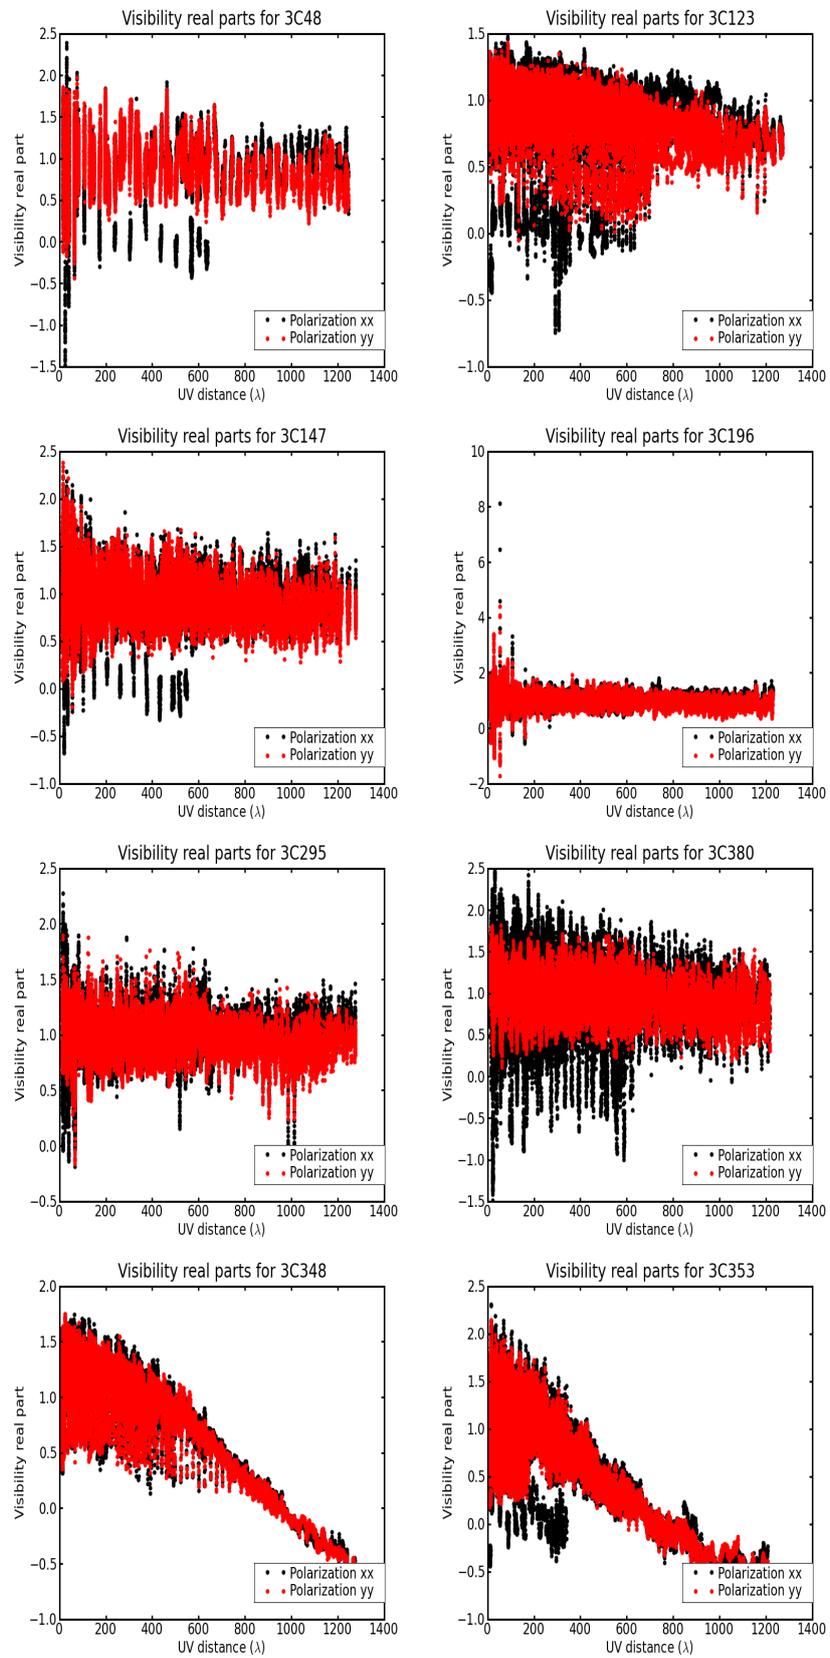
In the succeeding sections all the instrumental effects are calibrated. The order for calibrating depends on the importance of all the effects. The order is:

- Bandpass calibration. This affects the visibilities as function of frequency and has to be done just once.
- Ionospheric phase shift calibration. Has to be done before flux calibration because phase shifting repositions and focuses sources in the image plane.
- Flux calibration. If the fluxes and celestial positions change significantly during this step, phase shift and flux calibration should be repeated. Phase shift calibration improves the positioning and focusing of the sources in the image plane. As a result the fitted source fluxes increase during flux calibration afterwards.
- Gain calibration. This step is essential to tie the different observations to a common flux scale.

4.3 Local Sky Models for our fields

For doing self calibration we need a local sky model (LSM) for each field. In the measurement equation 4.15 it is the collection of coherencies \mathbf{B}_s and corresponding celestial positions (l_s, m_s) . An LSM can be very simple (one point source), more complicated (a collection of point sources), and very complicated (a collection of extended sources or even a 2D structure). Here, we have chosen LSMs of point sources, because modeling of extended sources was too time consuming for this project.

If we want to use a collection of point sources as LSM, at least the main source should be more or less a point source. To verify this, Fig. 4.2 shows the real parts of the visibilities as a function of uv distance for all the fields. If the main source dominates the field and is a point source, the real parts of the visibilities should be constant for all distances. The figure shows that sources 3C348, 3C353, Cas A, Cyg A, Tau A, and Vir A cannot be considered as point sources. We will not treat these sources in the remainder of this



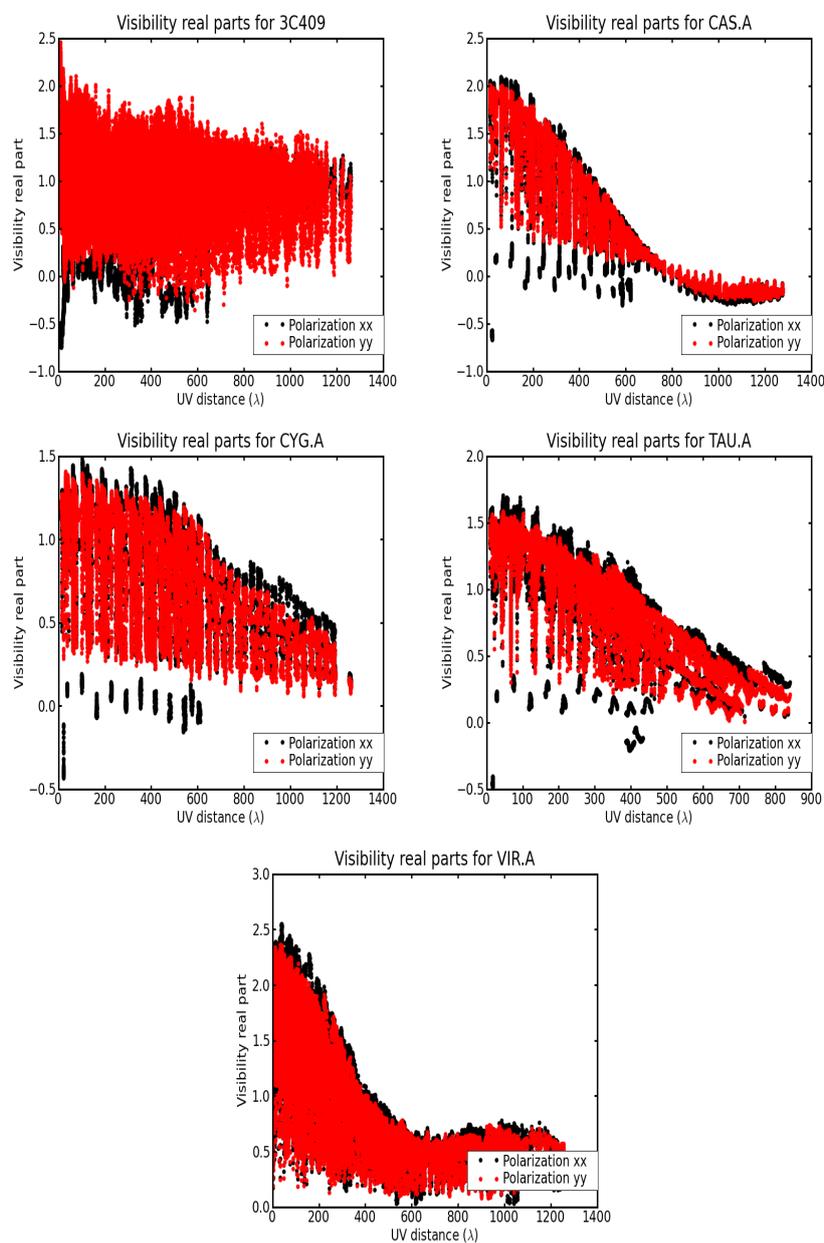


Figure 4.2: *Unflagged real parts of the visibilities for all the observed fields as a function of uv distance at 142 MHz for the fourth subband. Polarization xx is shown in black and polarization yy is shown in red.*

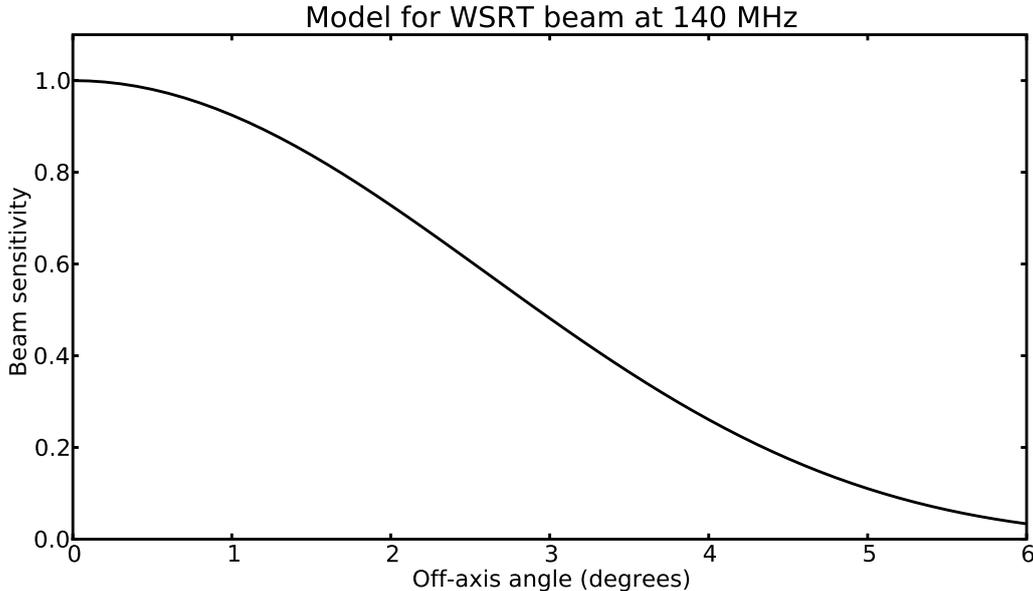


Figure 4.3: A model of the WSRT antenna beam, which is $\cos(1.2 \cdot 10^{-9} \nu d)^6$, where ν is the observed frequency in Hertz and d is the celestial distance from the pointing center in radians.

report. Unfortunately, by rejecting these sources, we lose all the primary calibrators of Sect. 1.1. Therefore fitting primary calibrators to literature values is not possible as well as transforming relative fluxes to absolute flux densities.

Figure 4.2 shows some deviating behavior for 3C123 and 3C196. 3C196 shows some visibilities with too big real parts and 3C123 shows many visibilities with too low real parts. These deviating visibilities are caused by instrumental effects: they disappear after doing self calibration as shown in Fig. 5.1.

The constant real parts of the visibilities as a function of uv distance in the remaining fields indicates that they are dominated by a very compact source.

To start self calibration we need an initial LSM for each field. This initial LSM should already contain the brightest sources in the field. During flux calibration the coherencies and celestial positions of the LSM are adapted, but no new sources are introduced.

An initial LSM can be obtained by taking the sources from an existing catalogue at a different frequency. At our frequency their fluxes might be wrong, but their celestial positions are likely correct. Here, we have obtained the initial LSMs from the Westerbork Northern Sky Survey (WENSS, for 3C48, 3C123, 3C147, 3C196, 3C295, 3C380) and the VLA Low-Frequency Sky Survey (VLSS, for 3C409 which is out of reach of the WENSS survey). The WENSS is a 325 MHz radio survey that covers the whole sky north of $+30^\circ$ declination to a limiting flux density of approximately 18 mJy (5σ) (Rengelink et al., 1997). The VLSS is a 74 MHz radio survey that covers the whole sky north of -30° declination.

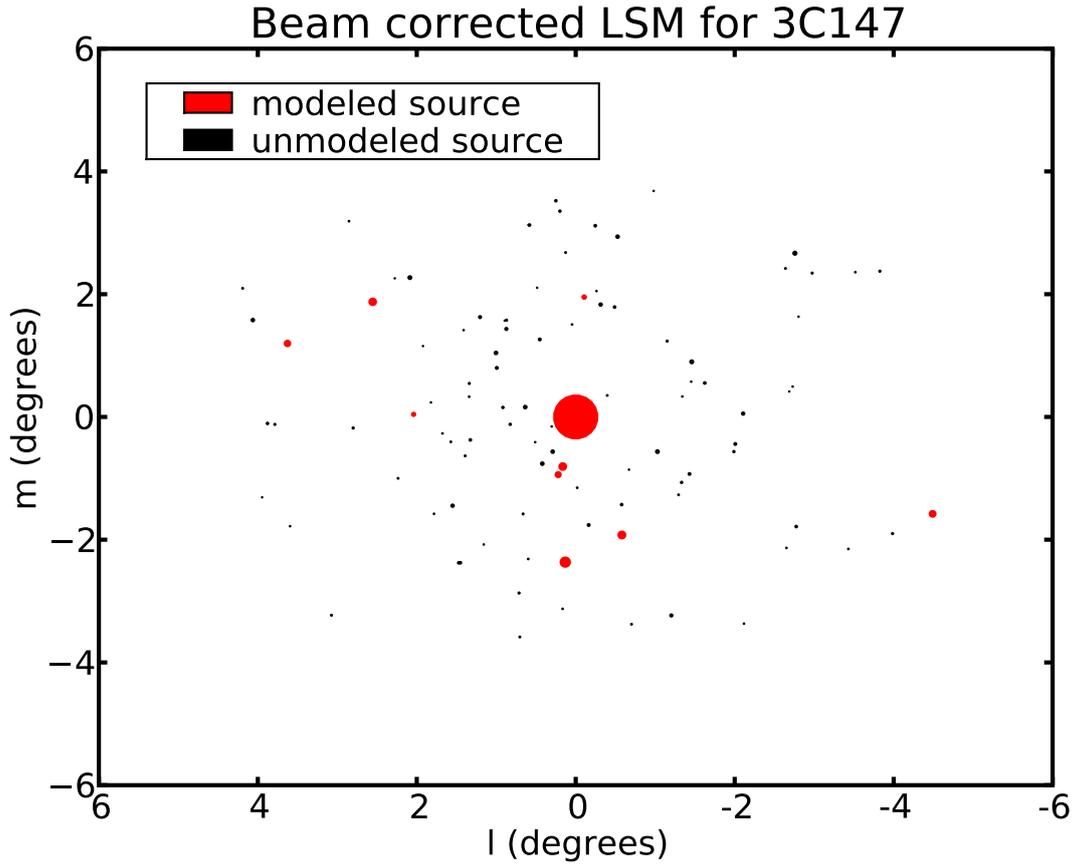


Figure 4.4: *Initial LSM for 3C147. Celestial positions and fluxes were taken from the WENSS survey. All fluxes were corrected using a beam model of $\cos(6.6 \cdot 10^{-8}\nu d)^6$, where ν is the observed frequency in Hertz and d is the celestial distance from the pointed direction in radians. The flux of a source is given by the area of the corresponding circle. The 10 brightest sources (which were calibrated for in Sect. 4.7) are shown in red. Right ascension and declination coordinates from the surveys were transformed to l and m with respect to the main source using Eq. (4.16). Initial LSMs for the other fields are shown in Fig. A.3*

This survey provides images with a resolution of $80''$ and with an average rms noise of 0.1 Jy/beam (Cohen et al., 2007).

The surveys contain unpolarized flux densities. However, we scaled the observations of each field to have a amplitude median of 1 (see Sect. 3.6). This should be done for the initial LSMs as well: relative normalized fluxes were included.

The number of sources in an LSM should be kept small and we included only ten sources, because flux calibration in calibration package MeqTrees (Noordam and Smirnov, 2009) is difficult (more on this in Sect. 4.7). The number of sources in a 12° field is typically much more in the WENSS survey. Therefore it is important to choose the right ones. Choosing the ten brightest sources from the survey is not enough: due to the beam of the antenna, off-axis sources appear dimmer. Better candidates are picked if the flux densities are first corrected for estimated beam gains. A good estimate for the beams of the WSRT is the function $\cos(1.2 \cdot 10^{-9} \nu d)^6$, where ν is the observed frequency in Hertz and d is the celestial distance from the pointing center in degrees. This beam is shown in Fig. 4.3. The figure shows why we have chosen 12° as field-of-view when selecting sources from the two surveys. The beam corrected LSMs are shown in Fig. 4.4 and Fig. A.3. Beam corrected flux densities are shown by the circle areas. The ten brightest sources that are picked for our LSMs are shown in red. The field 3C123 was just at the edge of the WENSS survey and therefore off-axis sources do not exist for all declinations. These figures confirm that the targets dominate the fields-of-view.

The celestial positions in the initial LSMs should be direction cosines (l, m) as described in Sect. 4.1. The right ascensions and declinations in the surveys were transformed to direction cosines using the sin geometry functions (Greisen, 1983):

$$l = \cos(\delta) \sin(\Delta\alpha), \quad (4.16a)$$

$$m = \sin(\delta) \cos(\delta_0) - \cos(\delta) \sin(\delta_0) \cos(\Delta\alpha), \quad (4.16b)$$

where $\Delta\alpha$ is the right ascension relative to the main source, δ is the declination of the source and δ_0 is the declination of the main source.

At this stage, we should have included the A-team sources (Cas A, Cyg A, Tau A, Vir A) in the LSM because these sources are far away, but very bright. Sections 4.9 and 4.10 will show the importance of putting these sources in the LSM. However, the A-team sources are all complicated sources and cannot be considered as point sources.

4.4 Setting up self calibration in MeqTrees

Self calibration was done using MeqTrees. MeqTrees (Noordam and Smirnov, 2009) is a software package for doing simulation and calibration of radio astronomical data. The package has been developed mainly by Jan Noordam and Oleg Smirnov at Astron in the Netherlands. Measurement equations are implemented in MeqTrees by defining data, transformations and relations using nodes in a hierarchy (called a tree). Implementation is done in a Python script and execution can be done from a graphical user interface as

well as from the command-line.

Setting up your own calibrations in a Python script for MeqTrees is not easily understood as of 2009. The beginners guide is in a preliminary state: it deals with the concepts of MeqTrees, but not on how to apply them to radio astronomy. On the other hand an example for doing calibrations in radio astronomy exists, but the given scripts use 'do-it-all' classes that are of no use when you need a different approach to self calibration.

Doing bandpass, phase, and gain calibrations was no problem with MeqTrees. However, doing flux calibration was difficult. We wanted to determine LSM fluxes and celestial positions that are not a function of time or frequency. Therefore all the data for one field should be used at once. Although our datasets for individuals fields were quite small (typically 2 hours of observing), when loading them into the system the system runs out of memory and solving becomes practically impossible. In Sect. 4.7 we solved this by calibrating fluxes for parts of the observations and combining the results afterwards.

Our python script for self calibration can be found in Sect. A.1. It is used in all the subsequent steps, each with different settings. For all the calibration steps the maximum number of iterations is 200 and the (u, v, w) values of antennas are calculated on-the-fly from their positions and the phase center.

4.5 Step 1: Bandpass calibration

The first step in calibration is bandpass calibration (the order is given at the end of Sect. 4.2). This calibration step models the bandpass filters of all antennas and polarizations of the WSRT.

During bandpass calibration amplitudes and phases are solved for each antenna, polarization, and frequency channel, limited to some constraints:

- The phases are limited to values between $-\pi$ and $+\pi$.
- The amplitudes should never be negative.
- The phases of both polarizations of antenna 0 were fixed at 0 at all frequencies. This constraint is necessary because the visibilities we used for calibration only contain phase differences between two antennas, not explicit phases of antennas. This antennas was never flagged.

A simple LSM of just one central source was used since the fluxes of the other nine sources in our initial LSMs are not yet calibrated.

Solving without time dependency is only allowed if the WSRT bandpass filters are stable. To test this stability within the four sessions, bandpass calibration was first tried separately on the 3C147 observations of each session. During the four sessions 3C147 was observed once, once, four times, and once. The bandpass solutions of the four sessions

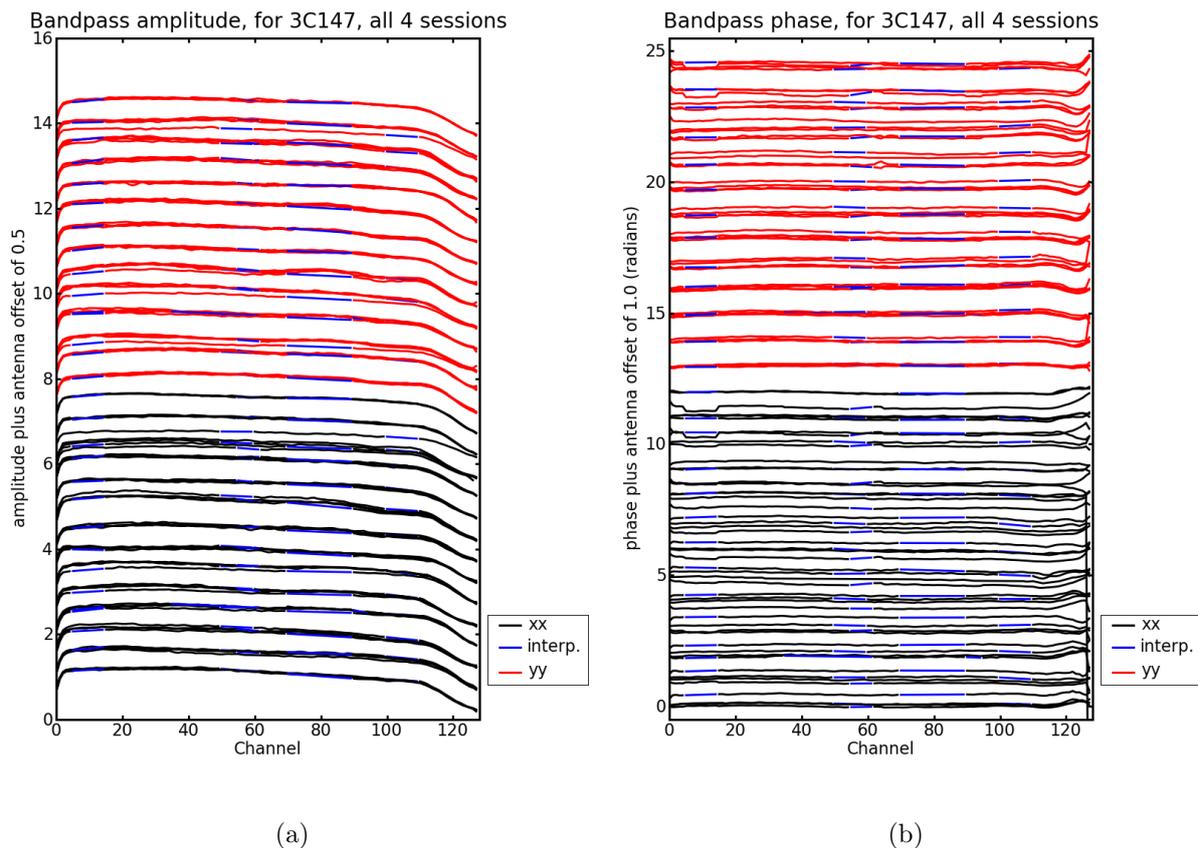


Figure 4.5: *3C147* bandpass solutions for amplitude (a) and phase (b) of the fourth subband for the four sessions. For the phases antenna 0 has no solutions, because this was one of the constraints. The offset increases per antenna, so the lower curve is of the first antenna and the upper curve is of the last antenna. The polarizations are placed on top of each other: Polarization *xx* is black and *yy* is red. The blue curves show interpolations on flagged channels. Antenna D was flagged during the first session, so only three curves are shown for this antenna.

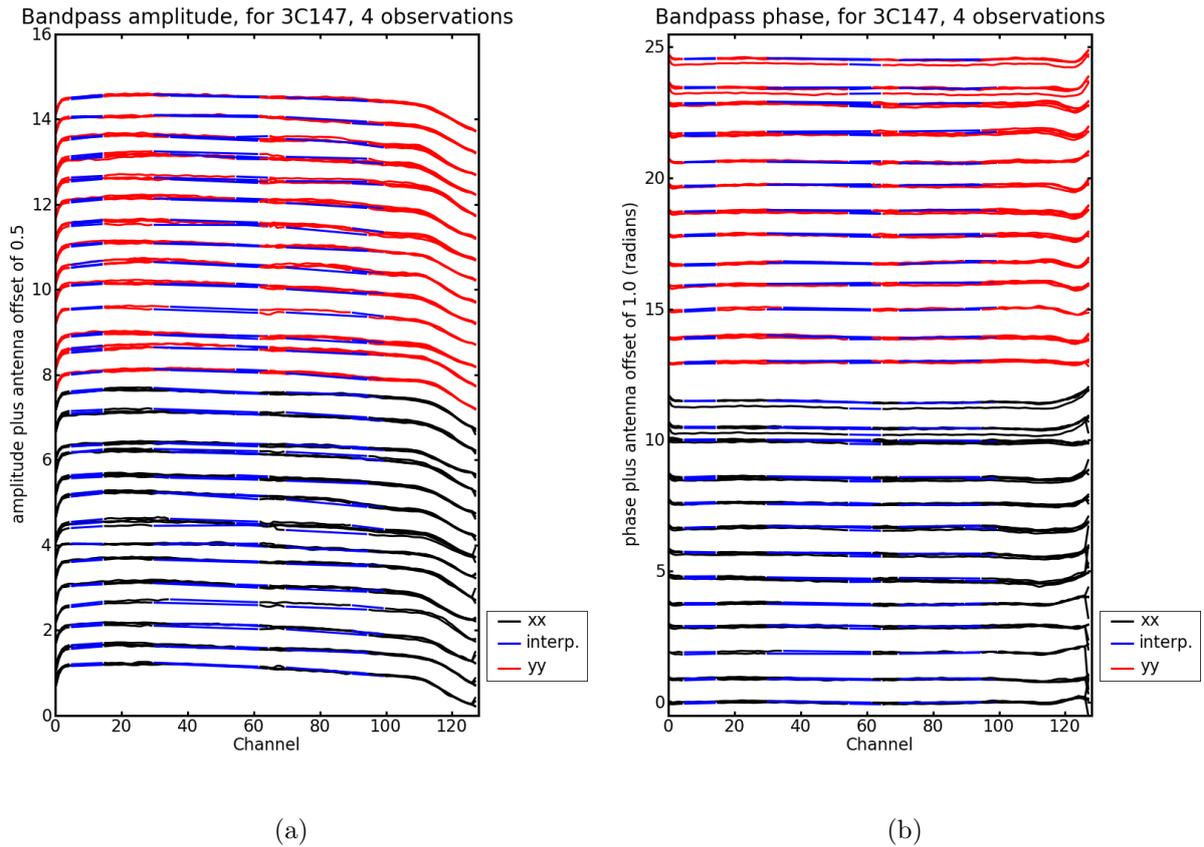


Figure 4.6: *3C147* bandpass solutions for amplitude (a) and phase (b) of the fourth subband for the four observations done during the third session. For the phases antenna 0 has no solutions, because this was one of the constraints. The offset increases per antenna, so the lower curve is of the first antenna and the upper curve is of the last antenna. The polarizations are placed on top of each other: Polarization *xx* is black and *yy* is red. The blue curves show interpolations on flagged channels. Antenna 3 was flagged during the first two observations of this session, so only two curves are shown for this antenna.

are shown in Fig. 4.5. The phases change dramatically: the median standard deviation of all the channels of all the bandpasses is 0.17 radians. The bandpass filters are therefore unstable during the four sessions.

To test the stability within one session, bandpass calibration was also tried separately on the four 3C147 observations within the third session. The bandpass solutions of the four observations are shown in Fig. 4.6. The phases are quite stable: the median standard deviation of all the channels of all the bandpasses is 0.04 radians. When considering amplitudes only the shapes of the curves in Fig. 4.6 are important. The overall amplitude during the observations will be calibrated during gain calibration. The shapes of the amplitude curves look constant. We considered the bandpass filters constant within one session.

We considered the characteristics of the bandpass filters as strictly internal, so independent of the observed field. Therefore if a bandpass is constant over a session, within this session we can use the bandpass solutions of one field for all the other observations. As a result, we used the four sets of 3C147 bandpass solutions for the different sessions as shown in Fig. 4.5. Antenna D was flagged during the first session. This is of no concern because it was flagged for all the observations (so for all fields) of this session. Therefore no calibration was done on this antenna during this session. Also, for each session no solutions were found for channels which were flagged. Interpolation (the blue curves) on these channels appears justified because the curves of both amplitude and phase are smooth.

Figure 4.5(a) shows that the bandpass amplitudes are very low at the edges of the band. This results in a lower signal to noise ratio. The phase solutions in Fig. 4.5(b) also behave strangely at the frequency edges. Therefore in all subsequent calibration steps, we only use channels 8 (141.45 MHz) until 120 (142.55 MHz) of the fourth subband. Because applying the Hanning filter makes each channel linearly dependent on its neighbors, only the even channels were used for further calibration and analysis.

The amplitudes of the curves in Fig. 4.5(a) look constant for all the antennas and polarizations, except antennas B (xx polarization) and C (both polarizations). During gain calibration in Sect. 4.9 frequency independent amplitude variations in time are solved for. From the small variations seen here for most antennas and polarizations, we expect small variations in gains as well.

The phases in Fig. 4.5(b) show some bumpiness between channels 5 and 15 for some antennas and polarizations. This is probably caused by flagging of the channels on many of the other antennas and polarizations, so little data was available for calibration. This bumpiness was ignored.

4.6 Step 2: Phase calibration

The next step in calibration is phase calibration. This calibration step models the phase shifts caused by the ionosphere. Flux and gain calibration is only useful after this step. For flux calibration in the next section all the data of one field is needed at once: all the observations and all the solutions for instrumental effects. Therefore starting phase cali-

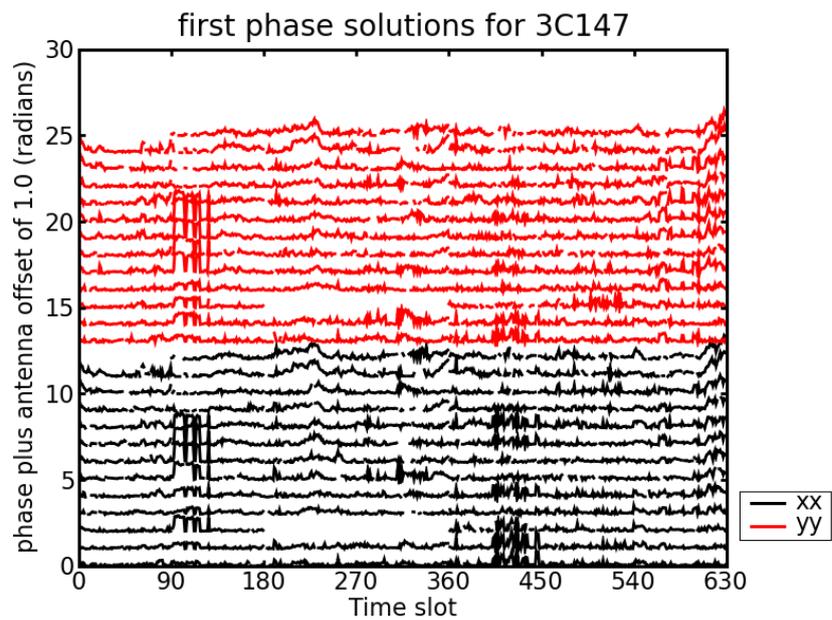


Figure 4.7: Phase solutions for 3C147. Phases are plotted with 1.0 offset per antenna. Antenna 0 has no solutions, because this was one of the constraints. The offset increases per antenna, so the lower curve is of antenna 1 and the upper curve is of antenna D. The polarizations are placed on top of each other: Polarization xx is black and yy is red. No solutions were found when the antenna was flagged during the time slot. The numbers of time slots are different per field because each field has its own number of used observations. Phase solutions for the other fields are shown in Fig. A.4

bration, the calibration is applied separately on each field: for each field, the observations were concatenated to one data set and fed to MeqTrees.

During phase calibration the phases are solved for each antenna, polarization, and time slot. The ionosphere changes the phases quickly and the changes depend on the location on the Earth’s surface. The phase shifts are considered constant within the field-of-view of the LSMs, which is 12° . Otherwise phases should be solved for different sub-fields separately. The constraints set on the phase parameters are:

- The phases are limited to values between $-\pi$ and $+\pi$.
- The phases of both polarizations of antenna 0 were taken 0 at all frequencies. This constraint is necessary because the visibilities we used for calibration only contain phase differences between two antennas, not explicit phases of antennas. This antenna was never flagged.

Again, a simple LSM of just one central source was used since the fluxes of the other nine sources in our initial LSMs are not yet calibrated.

All phase solutions are shown in Fig. 4.7 and Fig. A.4. The phase solutions show both slow and fast variations. The slow variations are easily followed due to the small 10 second integration times. For the slow variations often correlations are visible between antennas and polarizations at short baselines. This is expected since the ionospheric phase shifts depend on the location on the Earth’s surface and the phases are only expected to have a slope with increasing baseline within the size of the WSRT.

The fast variations on the other hand are too fast to follow. They typically have small durations and occasionally correlations are visible between antennas and polarizations. Although the fast variations look chaotic, we have not flagged time slots that are subject to fast variations. For some sources there are just too many (for example 3C380).

The phase solutions can be caused by the ionosphere, but also by shortcomings in our models. The observations show fringes from strong sources at high angular distances (probably the A-team sources) which are not part of the LSMs (more on this in Sect. 4.10). Phase calibration will try to compensate for the unmodeled sources which results in bad solutions.

4.7 Step 3: Flux calibration

The next step in calibration is flux calibration of each field. Flux calibration models the fluxes and celestial positions of the sources in the LSM. The initial LSMs are described in Sect. 4.3.

The sources are considered unpolarized. Therefore the coherency parameters are only the fluxes. During flux calibration the fluxes and celestial positions are meant to be solved once for all time slots and channels. The sources are considered stationary and to have

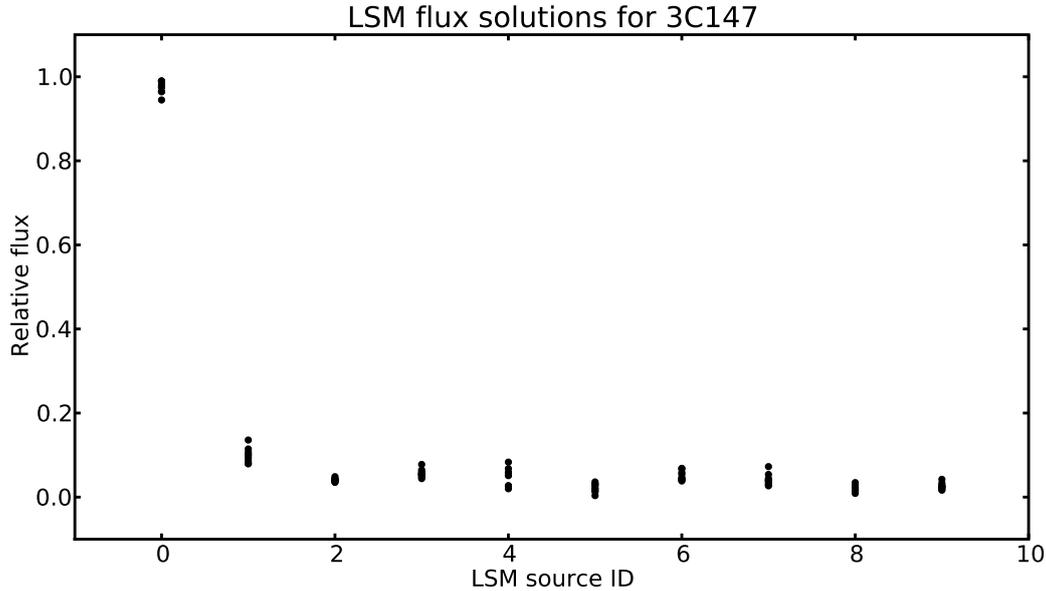


Figure 4.8: *LSM flux solutions for 3C147. There are nine solutions for each source, because the data set was divided into nine parts. The final fluxes for further modeling were obtained by taking the mean of the nine flux solutions for each source.*

no variability. The fluxes and celestial positions are also considered constant within our frequency subband (141.45 MHz to 142.55 MHz). The constraints set on the flux and celestial position parameters are:

- Fluxes should never become negative.
- The direction cosines l and m should stay between -1 and $+1$. It is unlikely that l and m become very far from zero since each initial LSM contains sources within 6° of the main source.
- The celestial position of the central source is not solved for. The central source is expected to dominate the field and therefore dominates the phase calibration. As a result, during phase calibration the central source is positioned in the phase center.

It is hard to do this calibration in MeqTrees. When all the observations for one field are loaded into the system, the system runs out of memory and solving becomes practically impossible. MeqTrees is capable of resampling the data, but this gave unexpected results. One way to avoid the problem is to cut all observations into slices, concatenate slices from different observations, and use the new datasets for calibration. This way the system does not run out of memory and each new data set contains data from all the hour angles. You need all hour angles so the sources become visible in the combined observations.

After flux calibrating this way a collection of solutions appear. An example of the flux solutions is shown in Fig. 4.8. Considerable variation in the found fluxes is visible for each

Table 4.1: *Flux solutions after calibrating for fluxes. For 3C409 there were only seven sources in the LSM.*

Source	3C48	3C123	3C147	3C196	3C295	3C380	3C409
1	0.892	0.910	0.975	1.035	1.035	0.894	0.674
2	0.128	0.027	0.100	0.087	0.041	0.154	0.073
3	0.047	0.015	0.041	0.080	0.032	0.066	0.052
4	0.047	0.021	0.057	0.107	0.032	0.065	0.418
5	0.051	0.010	0.049	0.031	0.032	0.038	0.036
6	0.029	0.013	0.022	0.057	0.048	0.031	0.027
7	0.083	0.044	0.051	0.031	0.019	0.019	0.023
8	0.029	0.013	0.042	0.047	0.020	0.029	–
9	0.034	0.012	0.021	0.043	0.020	0.034	–
10	0.032	0.016	0.026	0.028	0.021	0.028	–

source. The means of the flux solutions for the sources are used as final solutions for flux calibration. The same strategy is applied for the celestial positions of all the sources. The fluxes for all LSM sources of all the fields are shown in Tab. 4.1.

The flux solutions of the off axis sources are wrong (more on this in Sect. 4.10). This is probably caused by the assumption that the off axis sources appear unpolarized in the observations. However the instrumental polarization can be substantial within the 12° fields of the LSMs (up to 50 - 100%!). Our measurement equation does not deal with this instrumental polarization. This can be solved by introducing either beam calibration or polarization parameters for off axis sources during flux calibration. Unfortunately due to time constraints we were not able to do any of these actions.

4.8 Step 4: Repeat phase and flux calibration

The phases and fluxes were calibrated in the two previous calibration steps. However, the phases and fluxes can perhaps be improved. With the LSM from flux calibrating, phase calibration can be repeated. And if phases improve during this iteration, the sources are better focused. The LSM can then be improved as well by doing another flux calibration. And if the LSM has improved, the fluxes should become higher. If so, the complete process can even be repeated a number of times until the fluxes stop increasing.

This process was tested with all the fields. The new fluxes are printed in Tab. 4.2. Processing them typically took around 10 hours for one field. The new fluxes are not higher than the ones in Tab. 4.1 for all the fields, except for 3C123. Some are even lower. Therefore typically no improvements on the flux calibrations were made and no more iterations were performed.

Table 4.2: *Flux solutions after calibrating for fluxes a second time. The fluxes have not increased since the first flux calibration. For 3C409 there were only seven sources in the LSM.*

Source	3C48	3C123	3C147	3C196	3C295	3C380	3C409
1	0.881	0.945	0.973	1.037	1.036	0.897	0.616
2	0.130	0.028	0.100	0.088	0.041	0.151	0.083
3	0.058	0.010	0.042	0.081	0.032	0.065	0.037
4	0.055	0.023	0.053	0.108	0.033	0.067	0.347
5	0.056	0.009	0.050	0.018	0.028	0.040	0.094
6	0.023	0.011	0.017	0.060	0.049	0.023	0.080
7	0.089	0.010	0.052	0.033	0.010	0.028	0.050
8	0.026	0.012	0.038	0.049	0.011	0.026	–
9	0.042	0.007	0.034	0.039	0.019	0.042	–
10	0.034	0.007	0.028	0.014	0.019	0.034	–

4.9 Step 5: Gain calibration

The last step in calibration is gain calibration. Gain calibration models the sensitivity variations of all antennas and polarizations of the WSRT. The sensitivity variation is caused by gain variations of amplifiers, bandpass filters, and other components. The WSRT is a very stable instrument and the gain variations are expected to be only a few percent. The small amplitude variations found during bandpass calibration confirms this. This step was done after phase and flux calibration. Gain calibration is expected to improve our final flux estimation in Sect. 5.1 where we correct the data for all the modeled instrumental effects.

Gains can be solved for each antenna, polarization, and a certain time period. Gains are considered to have a strictly amplitude effect and the only solvable parameter is amplitude. If phase shifts are involved in the instrumental gains, they cannot be separated from the ionospheric phase shifts that also vary with antenna, polarization, and time slot. The amplitudes are considered frequency independent within our frequency subband (141.45 MHz to 142.55 MHz). Only one restriction was set on the amplitude parameters: they should never become negative.

Gain solutions for all the time slots in Fig. 4.9 and Fig. A.5 are much more variable than expected. They sometimes look correlated between antennas and this is not expected as well. The unexpected behavior is caused by imperfect LSMs: strong sources at large angular distances such as the A-team sources are not modeled, despite the importance of these sources (more on this in Sect. 4.10). As a result, gain parameters are adapted to make the model look like the observations as much as possible. Therefore the gain parameters are used to compensate for the fringes from strong unmodeled sources. The variation in gain solutions can be highly dependent on the polarization (like for 3C409). This is caused

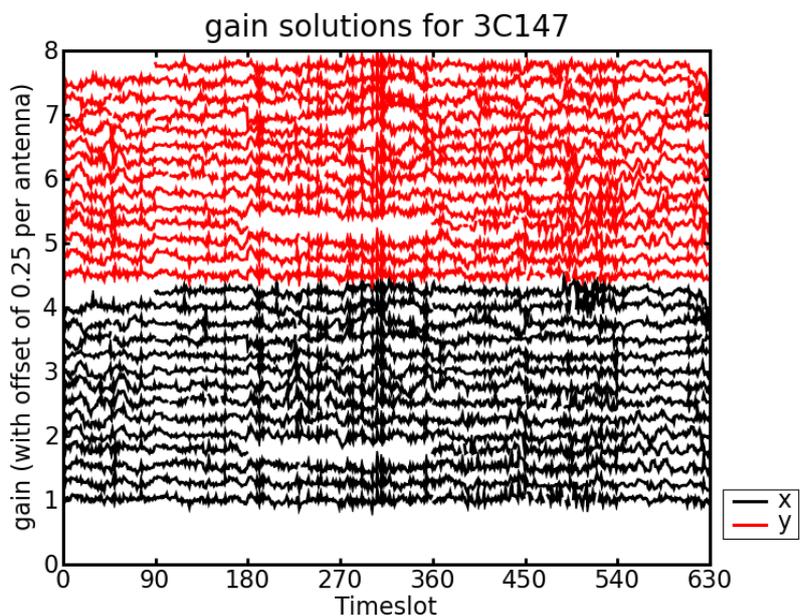


Figure 4.9: Gain solutions solutions for all time slots of 3C147. They are plotted with 0.25 offset per antenna. The offset increases per antenna, so the lower curve is of antenna 0 and the upper curve is of antenna D. The polarizations are placed on top of each other: Polarization xx is black and yy is red. No solutions were found when the antenna was flagged during the time slot. The numbers of time slots are different per field because each field has its own number of used observations. The gain solutions for the other fields are shown in Fig. A.5.

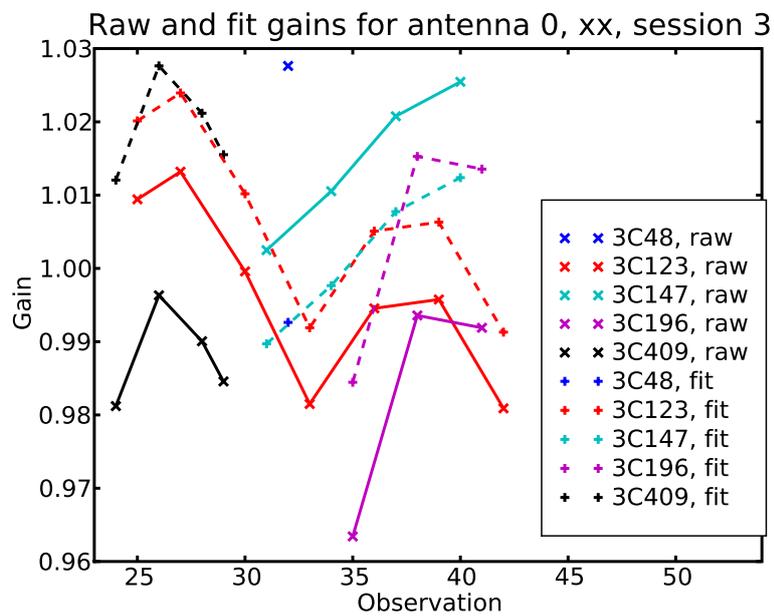


Figure 4.10: *Typical gain solutions solutions for the 12 observations of the seven sources during the third session. The diagonal crosses show the raw gain solutions and the upright crosses show the gain solutions after fitting the solutions of each field to one another. The corrected gain solutions of all fields together look more like a smooth curve than the uncorrected ones. Gain solutions are connected with lines for fields with multiple solutions within the session. Gain solutions for the other sessions are shown in Fig. A.8*

by a high instrumental polarization of the WSRT at large distances from the antenna axis. The ionospheric phase shifts will also be different away from the antenna axis. So again, it is important that the strong sources at large angular distances are in the LSM. And if sources at large angular distances are included in the LSM, instrumental polarization becomes an issue and beam calibration should be applied.

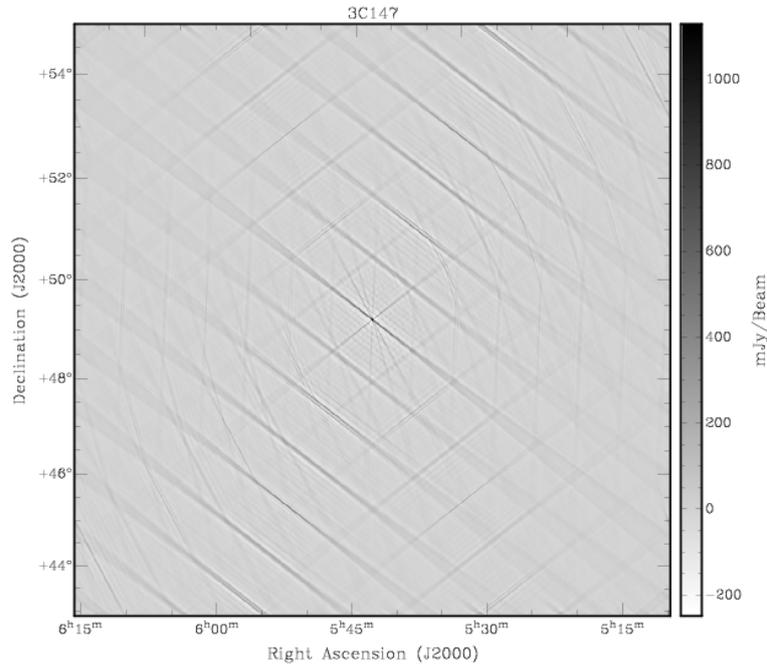
Calibrating gains for individual timeslots was not successful. Therefore the gain solutions were dropped and gains were subsequently calibrated for individual observations of 90 time slots. This way gain calibration can not compensate for the unmodeled strong sources. Gain solutions for a typical antenna and polarization and for the 12 observations of the seven sources during the first session are shown in Fig. 4.10, indicated by the diagonal crosses and the dashed curves.

The total flux in an observation is field dependent (despite the scaling in Sect. 3.6) and Fig. 4.10 shows that the gain solutions will be adapted to the observed flux. However the gains should be independent of the observed field: it should only show a fluently varying curve with time. Otherwise the relative fluxes cannot be obtained from the observations. The gain solutions are corrected to make them field independent. The gains as functions of time for all the fields are fit to one another. This was done by scaling the gain solutions of each field to the gain solutions of all the other fields and this procedure was performed three times. Scaling one function to another was done by deriving scaling factors from all the values of the first function and corresponding (interpolated) values of the other function and subsequently derive the right scaling factor using χ^2 minimization. The whole procedure was performed three times to damp the movements of all the functions. Typical corrected gain solutions are shown in Fig. 4.10 and Fig. A.8. Especially in Fig. A.8(c) the corrected gains look much smoother in time and independent of the fields. All the modeled gains were replaced by the new values and the 'corrected visibilities' were generated using the new gain solutions.

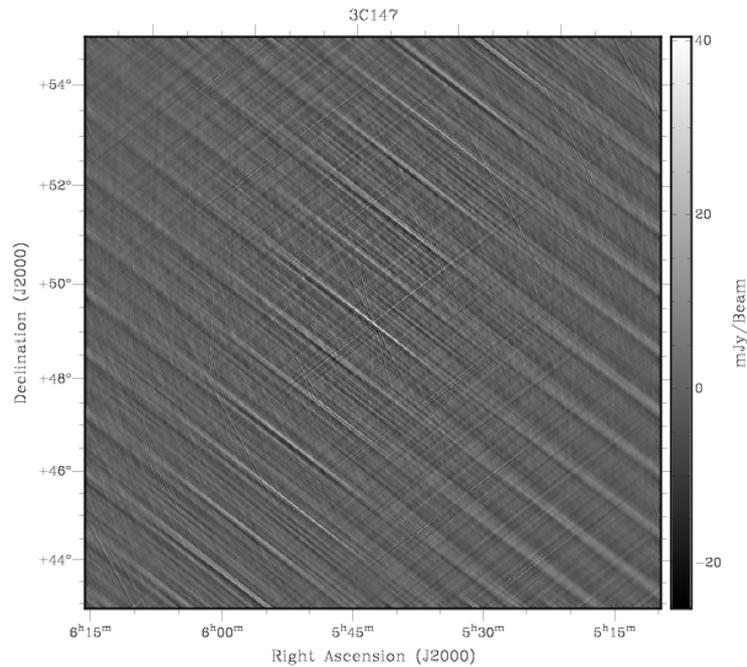
4.10 Self calibration quality

The quality of pre-processing and self calibration can be studied by inspecting residuals. A residual is created by subtracting the model from the observations. Afterwards this residual can be corrected for the instrumental calibration solutions. If the model perfectly describes the observations, the residual would just contain noise.

Inspecting residuals can be done in both the image plane and the visibility plane. Images of the observation and the residual of 3C147 are shown in Fig. 4.11. The contribution of off-axis sources is small relative to the main source because they are not visible in Fig. 4.11(a). The images were constructed using multi freq synthesis on frequencies 141.45 MHz to 142.55 MHz in 56 steps of 20 kHz. Radial weighting was used on the visibilities because this gives the best results for the WSRT. To obtain an image with a good PSF the uv plane must be filled uniformly. For the WSRT the uv density is higher for small



(a)



(b)

Figure 4.11: *Dirty images of the 3C147 field. (a) shows the image, uncorrected for modeled instrumental effects (negative gray scale, positive flux is shown in black). (b) shows the residual on a more sensitive grey scale (see colorbars), corrected for modeled instrumental effects. For both images multi freq synthesis was used on frequencies 141.45 MHz to 142.55 MHz in 56 steps of 20 kHz. Radial imaging weights were used because this gives the best results for the WSRT. The images span 12×12 degrees (0.7 arcmin/pixel). Dirty images of residuals of the other fields are shown in Fig .A.6*

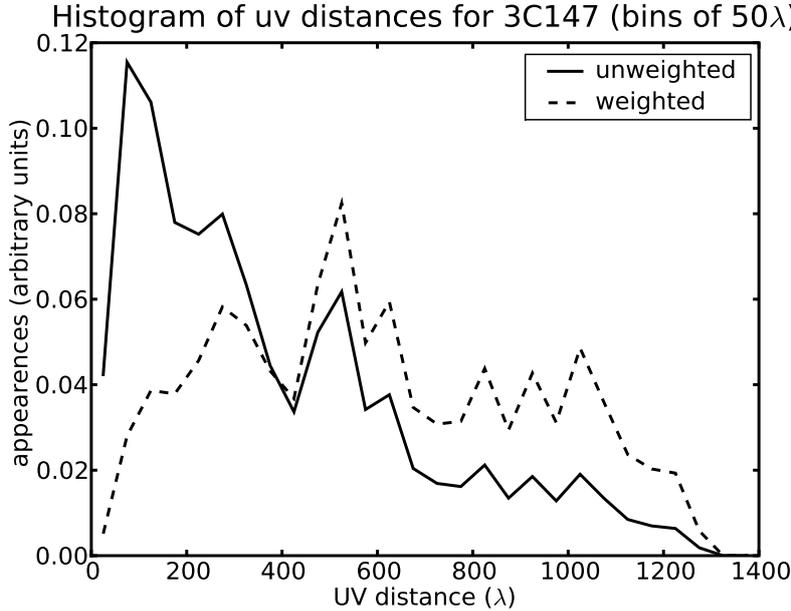


Figure 4.12: *Histogram of uv distances in 50λ bins at 142 MHz in all the used observations of 3C147 (solid curve). Radial weighting on the uv distances results in the dashed curve that is more flat.*

baselines as shown in Fig. 4.12 by the solid curve. After applying radial weighting the uv density becomes more flat as shown by the dashed curve in the figure. The images span 12×12 degrees (0.7 arcmin/pixel).

Self calibration did not succeed in successfully modeling the off-axis sources: the main source as well as off-axis sources are still visible in Fig. 4.11(b) and Fig .A.6. Strong patterns are also visible in these images. The patterns can be caused by strong unmodeled sources outside the field-of-view. The imperfections in calibrations have two causes: the absence of strong sources with large angular distances in the LSMs and the absence of instrumental polarization in the measurement equations.

The imperfect self calibration is also visible when inspecting residuals in the visibility plane. Visibilities of the observation and the residuals of 3C147 are shown in Fig. 4.13. Figure 4.13(a) and Fig. 4.13(b) have the same color scale. Figure 4.13(c) has a four times more sensitive color scale. The residual in Fig. 4.13(c) shows remaining fringes from unmodeled sources.

In Fig. 4.13(c) the remaining fringe has a phase delay for different frequencies. From this delay a rough guess can be made on the celestial distance of the source of the fringe. When looking at one time slot the phase changes by roughly 1 radian over the subband. Because the subband has a width of 1.2 MHz, 1 radian phase difference corresponds to $\pi/2\pi\Delta\nu = 0.4 \mu\text{s}$, which corresponds to 120m. Because the baseline between antennas 1 and 2 is 144 m at maximum, 120 m can correspond to a celestial distance of $\sin^{-1}(120/144) = 56^\circ$

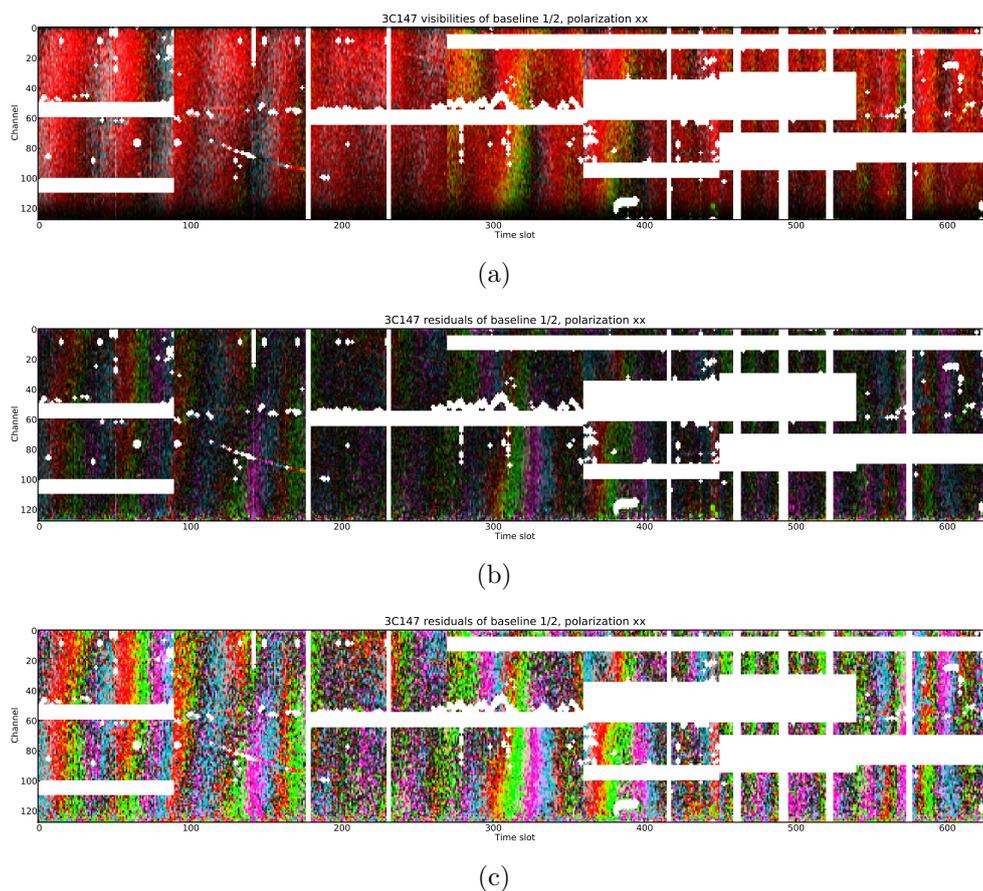


Figure 4.13: (a) and (b) show the visibilities and the residuals on the same color scale (an amplitude of 2 corresponds to the highest intensity) for the seven observations of 3C147 and for the fourth subband, baseline 1/2 and polarization xx . (c) shows the residuals on a four times more color scale. (a) is not corrected for modeled instrumental effects. (b) and (c) are corrected for modeled instrumental effects. The color palette of appendix A.2 is used for displaying complex values. Residuals of the other fields are shown in Fig .A.7

(for smaller baselines the distance is even larger). This distance is similar to the distance to Cas A. Therefore the remaining fringe is caused by Cas A. Including the four A-team sources in the LSM is important at the 140 MHz regime. However, as mentioned before, the A-team sources are all complicated sources and were not included in our LSMs of point sources.

As visible in Fig. 4.13(c), the remaining fringes dominate the residual data. The thermal noise has a much lower amplitude.

For the six other fields residuals in the image plane are shown in Fig. A.6 and their corresponding residuals in the visibility plane are shown in Fig. A.7. The image plane residuals show that calibration of flux and celestial position of the main source is not perfect for many fields. Sometimes there is a misalignment between the observation and model. This makes the main source look 'black and white' (3C196 and 3C295). This can be due to imperfect phase calibration causing bad focusing of the sources. Because of time constraints, no solutions for these problems (like flagging phase solutions at time slots) are presented.

Chapter 5

Data analysis

5.1 Obtaining relative fluxes

Table 5.1: *The radial weighted means of all fields and corresponding errors are shown in columns two and three. Columns four and five show the means and errors after inverting the scaling of Sect. 3.6. They are the obtained relative fluxes.*

Field	Weighted mean	Error	Unscaled mean	Error
3C48	0.9589	0.0016	0.002850	0.000005
3C123	0.9441	0.0009	0.010518	0.000010
3C147	1.0080	0.0012	0.002904	0.000003
3C196	1.0194	0.0017	0.003486	0.000006
3C295	1.0197	0.0017	0.003735	0.000006
3C380	1.0041	0.0013	0.003442	0.000005
3C409	0.6366	0.0006	0.002952	0.000003

For all the fields we want to obtain the (unitless) flux of the main source. These fluxes can be obtained from the observations and the calibration models. For each field the contribution of the modeled off-axis sources can be subtracted from the visibilities and the remaining visibilities can be corrected for all the modeled instrumental effects. The remaining visibilities, called the 'corrected visibilities', contain the flux of the main source, noise, and unwanted flux from badly or unmodeled off-axis sources.

We want to obtain the flux ratios of the main sources from the 'corrected visibilities'. The main source always has a completely real contribution to the visibilities, because after phase calibration the main source is always located in the phase center. Visibilities of off-axis sources run around the origin in the complex plane with time. By time averaging the real parts of the visibilities, the contributions by both off-axis sources and noise vanish while the main source remains. The off-axis sources can also be extended emission, for example from the galactic background. The flux of the main source can then be estimated

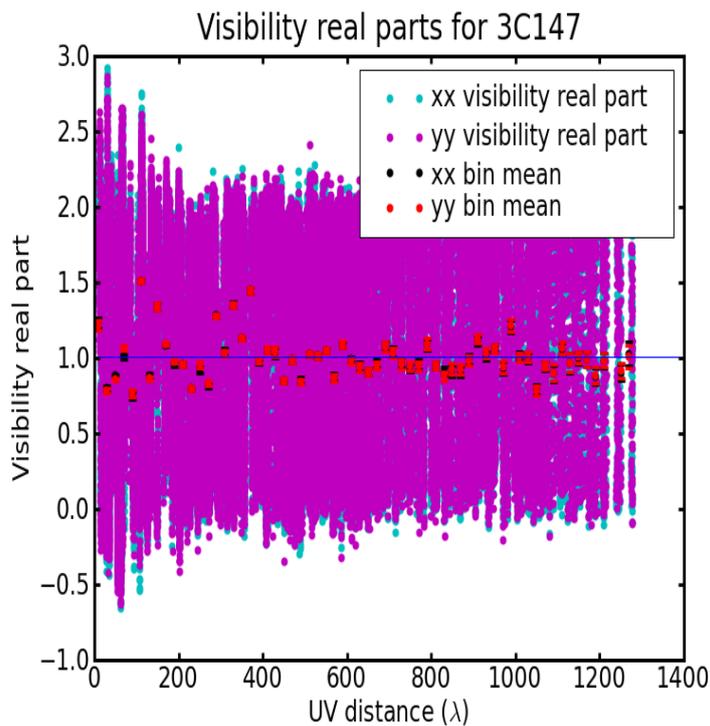


Figure 5.1: *Real parts of the 'corrected visibilities' (modeled off-axis fluxes subtracted and corrected for modeled instrumental effects) of 3C147 as a function of uv distance. Real parts were obtained by taking the mean of the real parts of channels 8-120. Cyan is the xx and magenta is the yy polarization. The black and red errorbars are means over 20λ bins for the xx and yy polarization, respectively. The blue horizontal line shows the radial weighted mean of the real parts. The real parts of the 'corrected visibilities' for other fields are shown in Fig. A.9*

by time averaging the real parts of the 'corrected visibilities' as described above. And this is independent of frequency and baseline. Therefore if the mean is taken over all the real parts of the visibilities of one field (over all the observation dimensions time, frequency, and baseline) the flux of the main source can still be found, despite the imperfect self calibration.

In Sect. 4.10 was shown that the best image quality was found when radial weighting was applied to the visibilities. This strategy can be applied here as well. The mean real part of the visibilities can be determined from weighted visibilities having a better PSF.

Before taking the fluxes from the real parts of the visibilities, we take a closer look at these real parts. The dependency of the real parts of the 'corrected visibilities' on uv distance is shown in Fig. 5.1 and Fig. A.9. For all fields a lot of variation is visible in the real parts for all the distances. Figure 4.13(c) already showed that these variations were caused by systematic errors due to remaining fringes and were not caused by statistic variation.

Figure 5.1 and Fig. A.9 can be compared to Fig. 4.2. In Fig. 5.1 the modeled off-axis sources are subtracted and all the data is corrected for the modeled instrumental effects. These corrections have a big effect on the data. However, the real parts in Fig. 5.1 still have a lot of variation, caused by the unmodeled strong sources.

To study if the systematic errors caused by the fringes depend on the baselines and polarizations we binned all the real parts of the visibilities with bin widths of 20λ . The means of the bins are shown as the black and red dots in Fig. 5.1 and Fig. A.9 together with their small errorbars (xx is black and yy is red). The bins show that systematic errors can be quite different for different baselines, but is very similar for the two polarizations. Because the variation on the bin means looks symmetric, taking the mean over all real parts for obtaining the flux of the main source can be justified. The polarizations similarity shows that the errors are really systematic and are caused by unpolarized source(s). If the sources were polarized the real parts would be different for the two polarizations.

The bin means can be considered constant with uv distance for all fields. There is no increase in amplitude at short baselines. Therefore it is fair to consider our studied sources as point like sources. No faint extended emission is visible at very short baselines (the shortest distance between antennas 9 and A are only 36 m or 17 wavelengths, see also Fig. 3.12).

The determined means of the real parts of 'corrected visibilities' using radial weighting for all fields are given in Tab. 5.1, in the second and third column. The errors in the weighted means are very small, because there are many visibilities in the observations. The weighted means were also plot in Fig. 5.1 and Fig. A.9 as the blue horizontal lines.

The relative fluxes of all the main sources can now be derived by taking the weighted means and invert the scaling applied to the data in Sect. 3.6. The relative fluxes are given in the last two columns of Tab. 5.1.

Precision while scaling relative fluxes to absolute flux densities is very important. Un-

fortunately we have dropped all primary calibrators from our observations. Therefore we cannot transform the flux ratios to absolute flux densities and cannot contribute in improving the absolute flux scale at low frequencies.

Chapter 6

Conclusions and suggestions for future work

The following conclusions were made during pre-processing:

- If absolute system temperatures are not required, it is better to use P_{off} instead of T_{sys} , because P_{off} is more representative of the average conditions during the observation.
- Flagging visibilities should be done after tapering and AGC compensation.
- Although unprotected and heavily used, some low frequency bands are not dominated by RFI and can be used for astronomical observations.
- Activity of the Sun can destroy your low frequency observation.
- Manually flagging frequency channels remains necessary, but is feasible because bad channels appear simultaneously for most baselines.

The following conclusions were made during self calibration:

- Recording cross-hand polarizations is necessary for leakage calibration.
- Creating an LSM of only point-sources for low-frequency self calibration is not sufficient, because the A-team sources should be included that are not point-sources. Including the A-team in the LSM is critical for doing good self calibration on low-frequency observations.
- The bandpass filters of the WSRT can be considered constant within a 12h observing session.
- Ionospheric phase variations are fast at low frequencies, but can be followed with 10 sec integration times.
- Flagging bad calibration solutions is critical in properly modeling observations. Residuals will otherwise show the imperfections.

The following conclusions were made during data analysis:

- Even with badly modeled observations, the relative fluxes of central sources can be obtained by time- and frequency-averaging.
- Sources with predictable fluxes should be observed to get a good scaling factor.

We suggest to redo the observations on the WSRT/LFFE, this time with full linear polimetry. The absence of the xy and yx correlation products is a big drawback in the existing observations, because strong sources at high celestial distances are detected and appear polarized. By including the xy and yx terms, the instrumental polarization can be modeled.

However, the highest possible spectral and time resolution should be used for the new observations. This way the loss of data due to temporary RFI acting on narrow frequency bands is kept as low as possible.

For the new observations we propose to focus more on radio sources with constant spectral indices at low frequencies, like 3C43, 3C89, 3C180, 3C196.1, 3C270, 3C272.1, and 3C441 (Kellermann et al., 1969). Getting the right scaling factor is critical in obtaining precise absolute flux densities.

Chapter 7

Acknowledgments

I would like to thank my supervisors Michiel Brentjens and Ger de Bruyn. It has been a very interesting project and it has been fascinating for me to experience how much labour is put into getting some results from a radio observation. I would also like to thank the various people at Astron/RUG that helped me with my first steps in radio astronomy and self calibration, especially naming Sarod Yatawatta, Jan Noordam, and Oleg Smirnov.

I would also like to thank all the astronomy students I have met the last four years. It has been a fun time with all of you and it has been great to witness how all of you combine fun with determination.

But most of all I would like to thank my girlfriend Paulien. Without her this master thesis would never have existed. She has been my big support the last four years.



Figure 7.1: Jeffrey (1) and Michiel (2, hardly visible) organizing an astronomical youth camp back in 1996. Some of the youngsters are now working in professional astronomy (3 = Steven Rieder, 4 = Ernst de Mooij, now both PhD students in Leiden).

Appendix A

A.1 Scripts

The Python script for data reduction and figure generation is available at:

http://www.astro.rug.nl/~bout/master_thesis/code_bout_master_thesis.py

The Python script for calibration is available at:

http://www.astro.rug.nl/~bout/master_thesis/bout_meqtrees_com_line.py
(use from the command-line)

http://www.astro.rug.nl/~bout/master_thesis/bout_meqtrees_gui.py
(use from the MeqBrowser GUI)

A.2 Color palette for displaying complex values

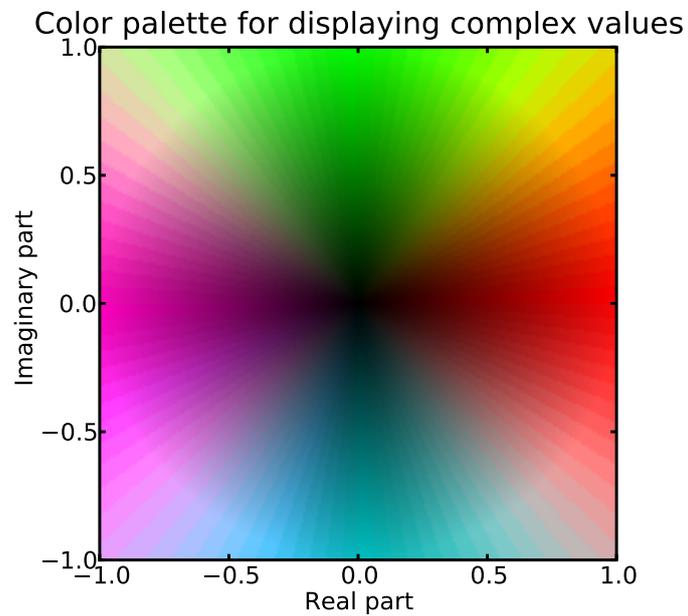
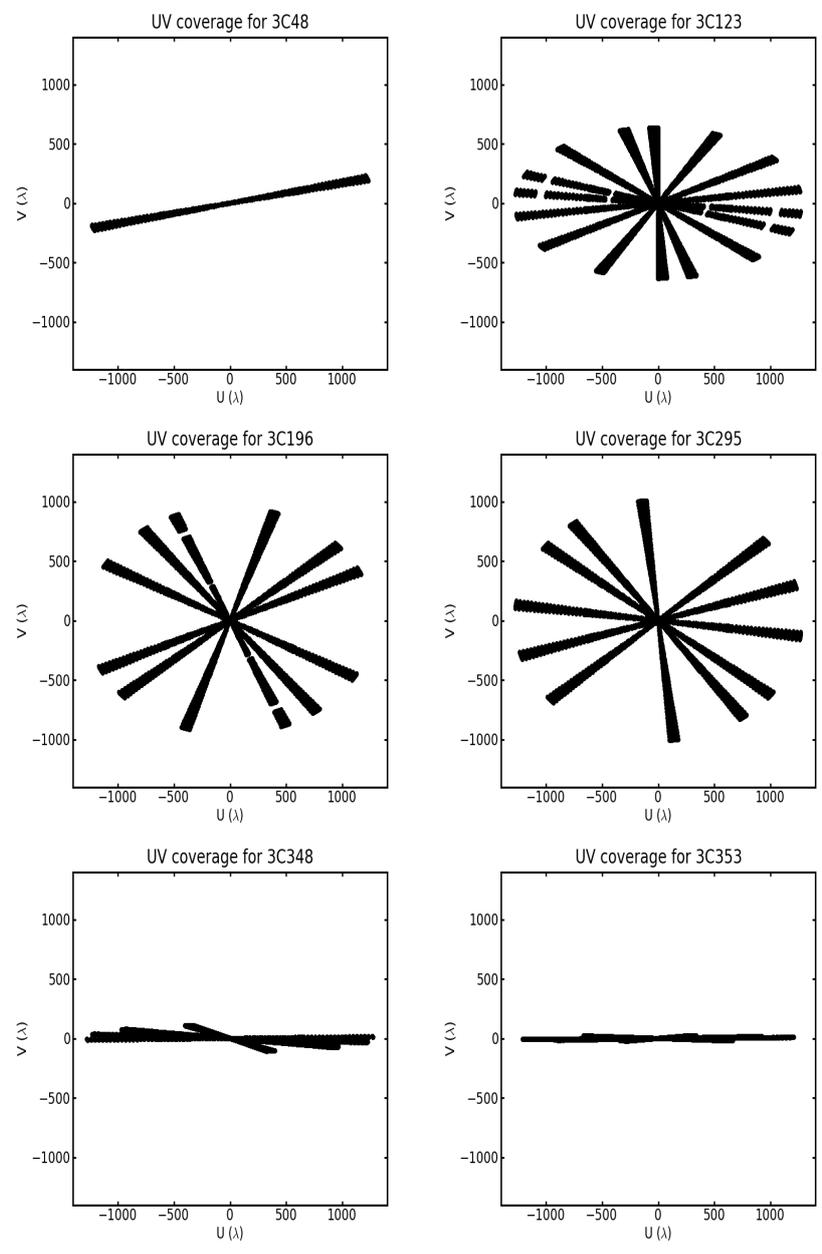


Figure A.1: *Color palette for displaying complex values as developed by Michiel Brentjens. With increasing angle the color changes from red to green, magenta, cyan, and back to red. The intensity increases with increasing amplitude. For amplitudes above 1, the maximum intensity is applied. In visibility figures using this palette, flagged visibilities are shown in white.*

A.3 Figures



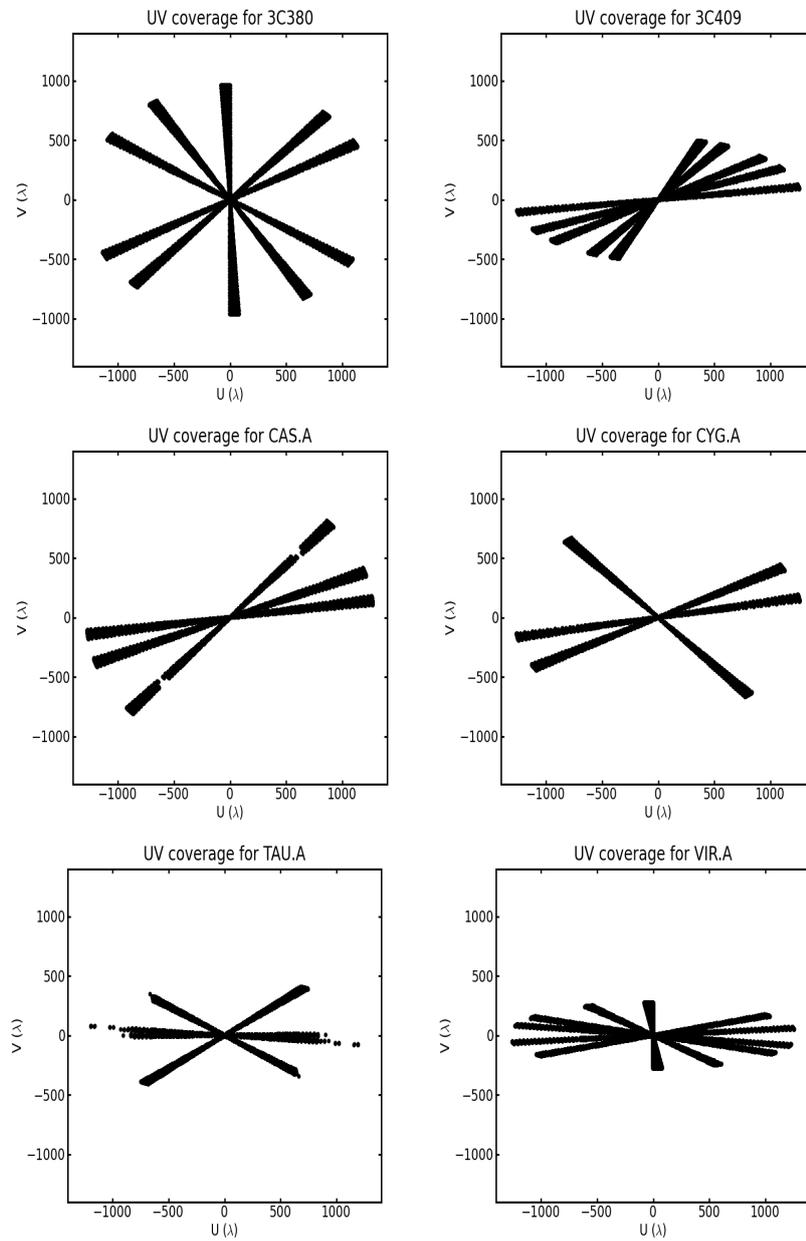


Figure A.2: The uv coverage at 142 MHz of all the observed fields (except 3C147) for the four sessions with usable data (as concluded in Sect. 3.3). The frequency 142 MHz is the one selected for processing in Sect. 3.3. The uv coverage of 3C147 is shown in Fig. 2.1

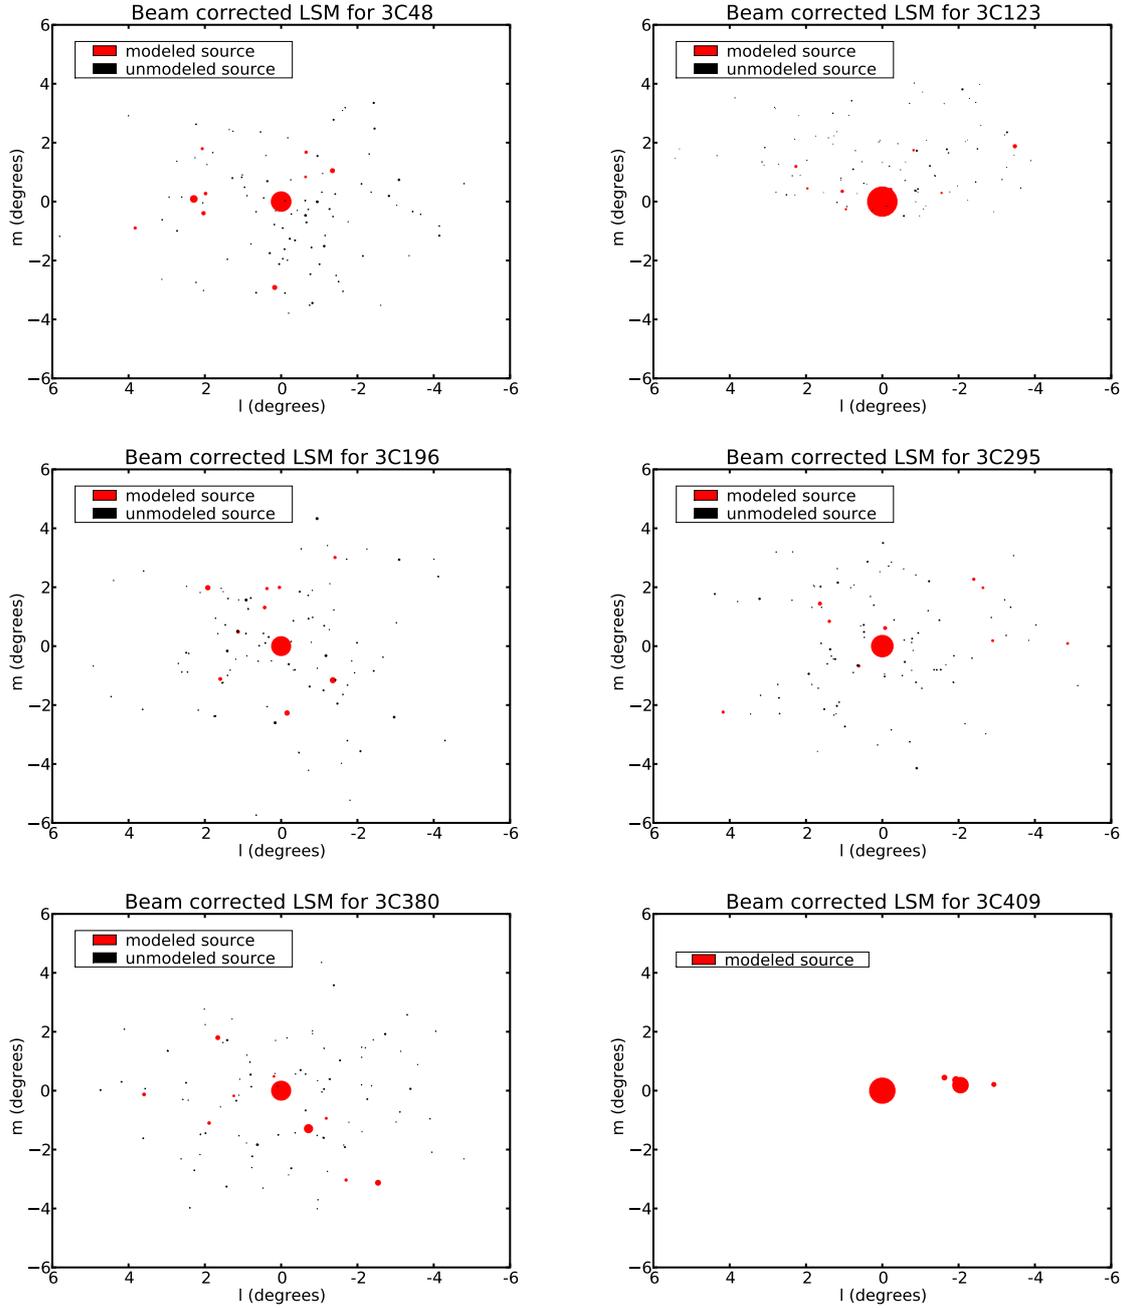


Figure A.3: Initial LSMs for the modeled fields (except 3C147). Celestial positions and fluxes were taken from the WENSS survey (all fields except 3C409) and the VLSS survey (3C409). All fluxes were corrected using a beam model of $\cos(1.2 \cdot 10^{-9} \nu d)^6$, where ν is the observed frequency in Hertz and d is the celestial distance from the pointed direction in degrees. The flux of a source is given by the area of the corresponding circle. The 10 brightest sources (which were calibrated for in Sect. 4.7) are shown in red. Right ascension and declination coordinates from the surveys were transformed to l and m with respect to the main source using Eq. (4.16). The initial LSM for 3C147 is shown in Fig. 4.4

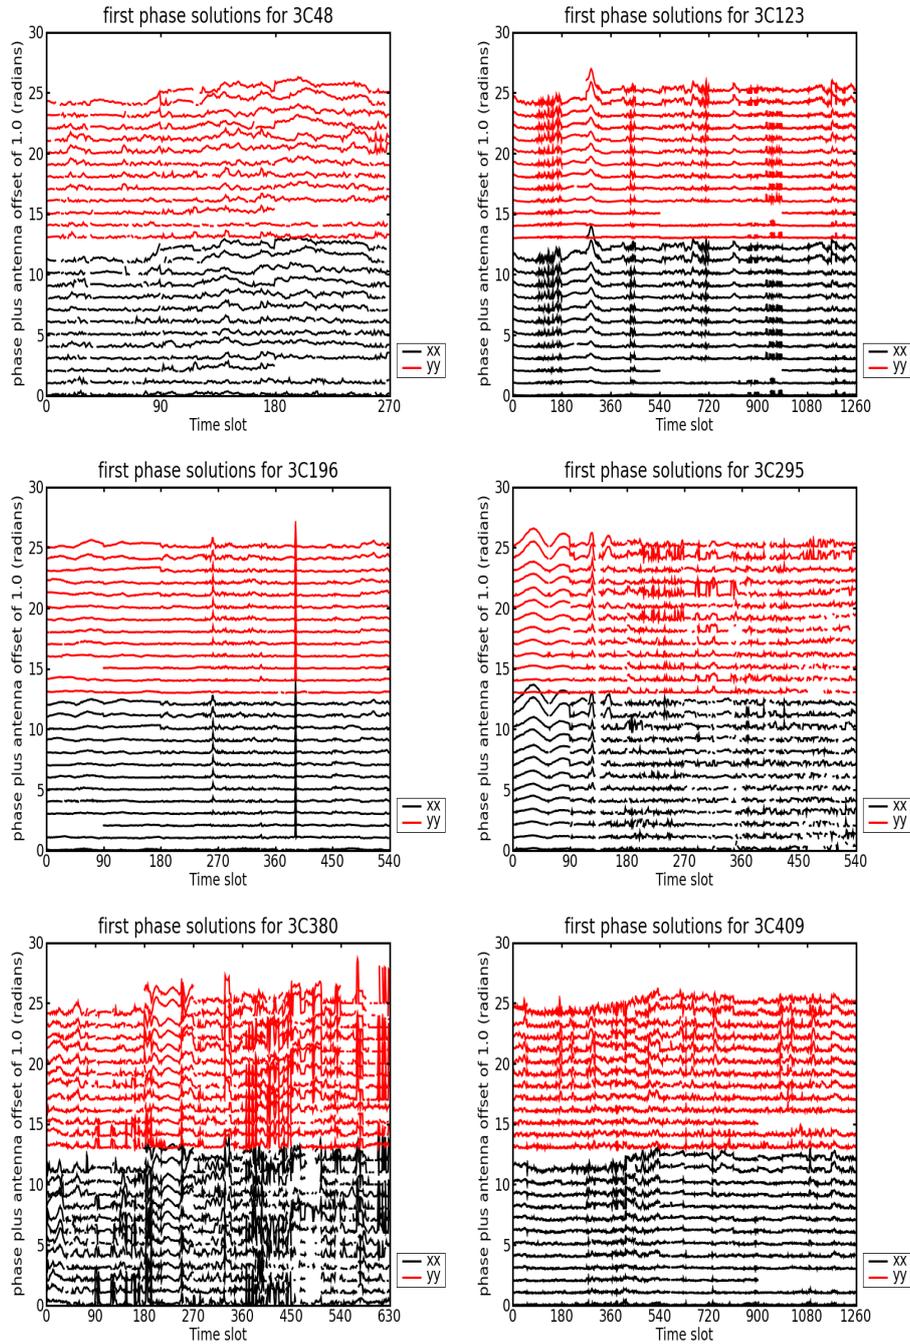


Figure A.4: Phase solutions for all the fields (except 3C147). Phases are plotted with 1.0 offset per antenna. Antenna 0 has no solutions, because this was one of the constraints. The offset increases per antenna, so the lower curve is of antenna 1 and the upper curve is of antenna D . The polarizations are placed on top of each other: Polarization xx is black and yy is red. No solutions were found when the antenna was flagged during the time slot. The numbers of time slots are different per field because each field has its own number of used observations. Phase solutions for 3C147 are shown in Fig. 4.7

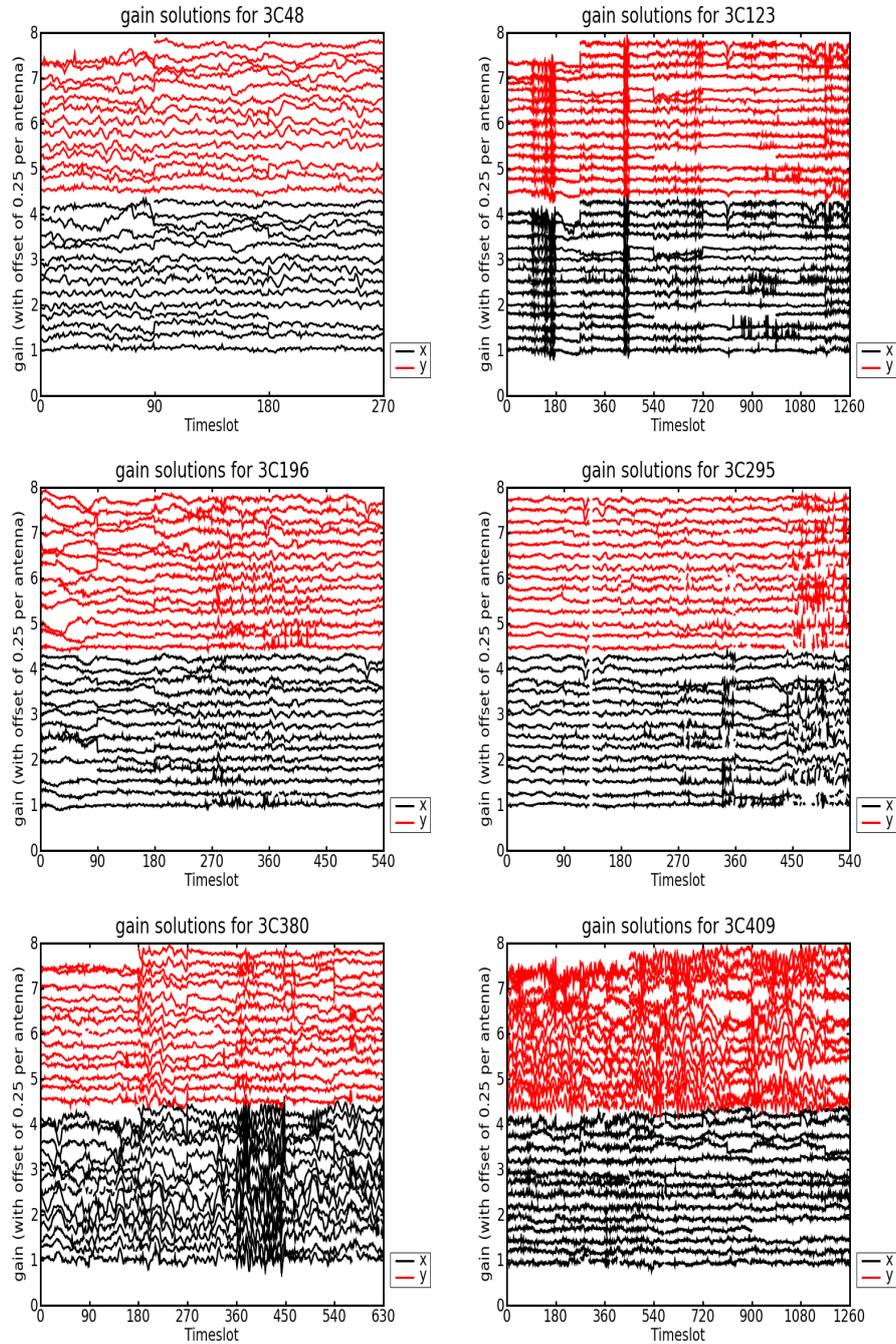
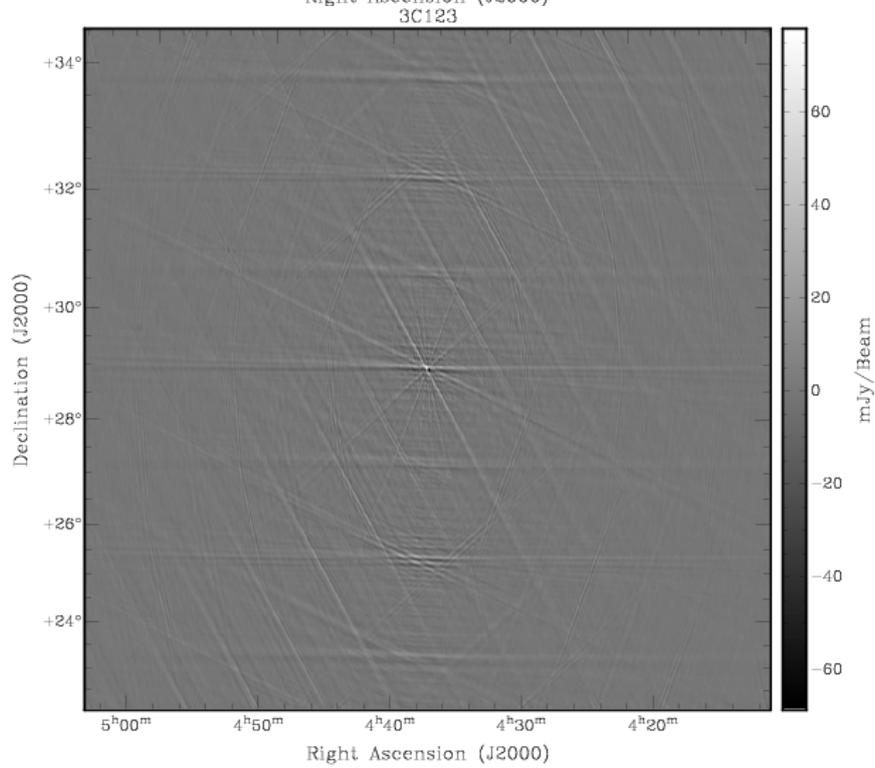
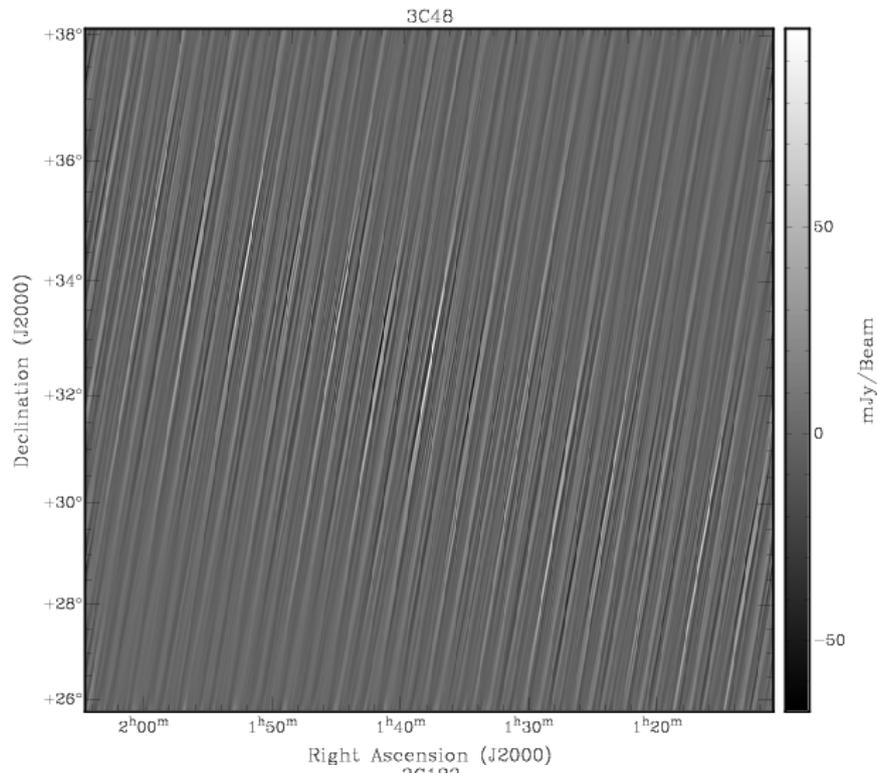
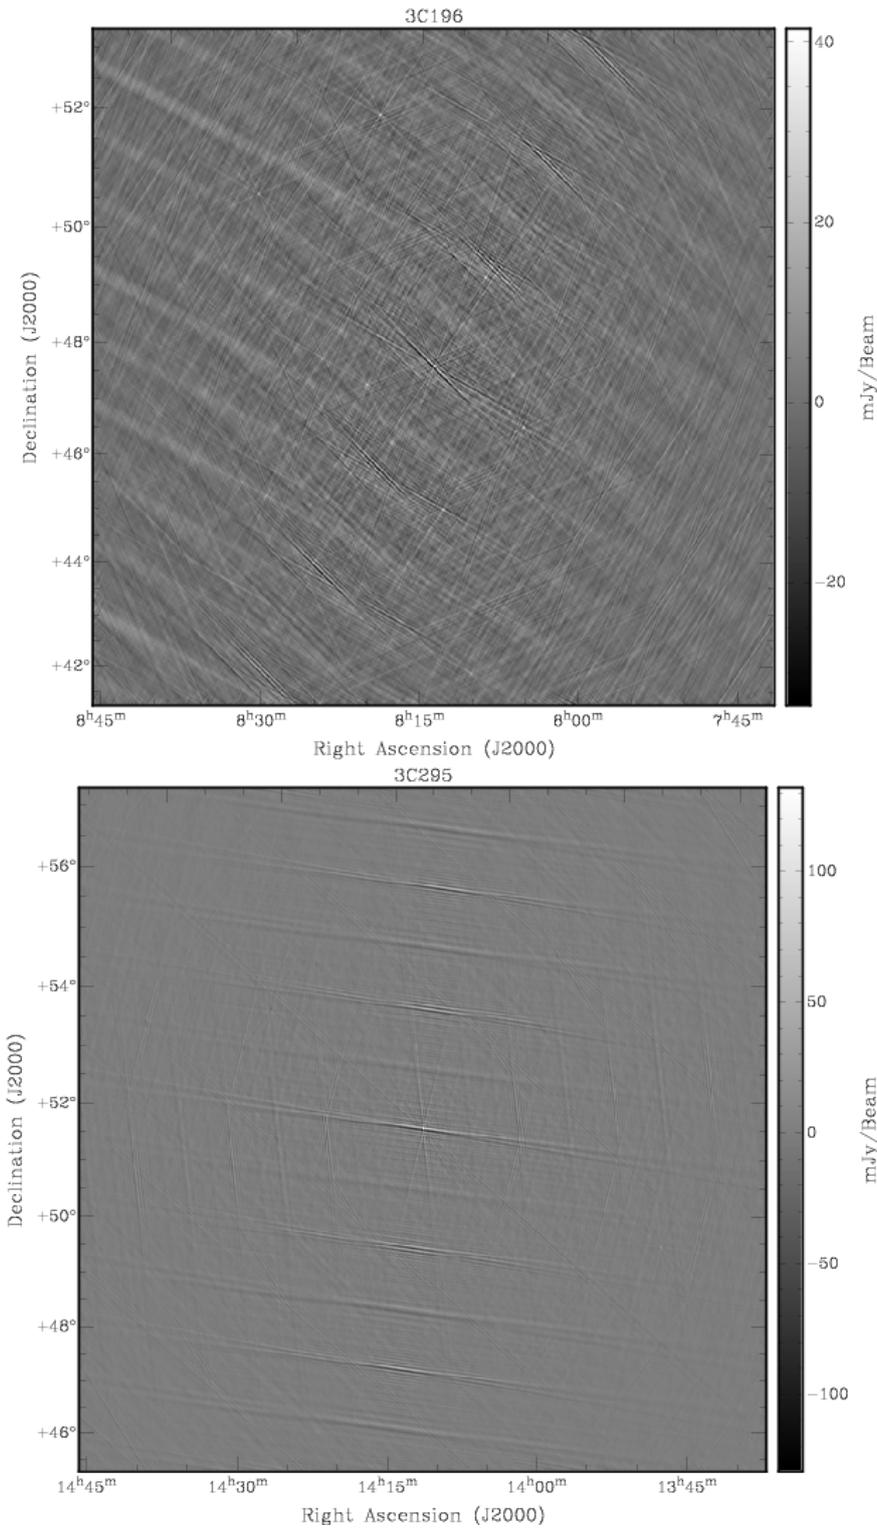


Figure A.5: Gain solutions solutions for the seven fields (except 3C147). They are plotted with 0.25 offset per antenna. The offset increases per antenna, so the lower curve is of antenna 0 and the upper curve is of antenna D. The polarizations are placed on top of each other: Polarization xx is black and yy is red. No solutions were found when the antenna was flagged during the time slot. The numbers of time slots are different per field because each field has its own number of used observations. The gain solutions for 3C147 are shown in Fig. 4.9.





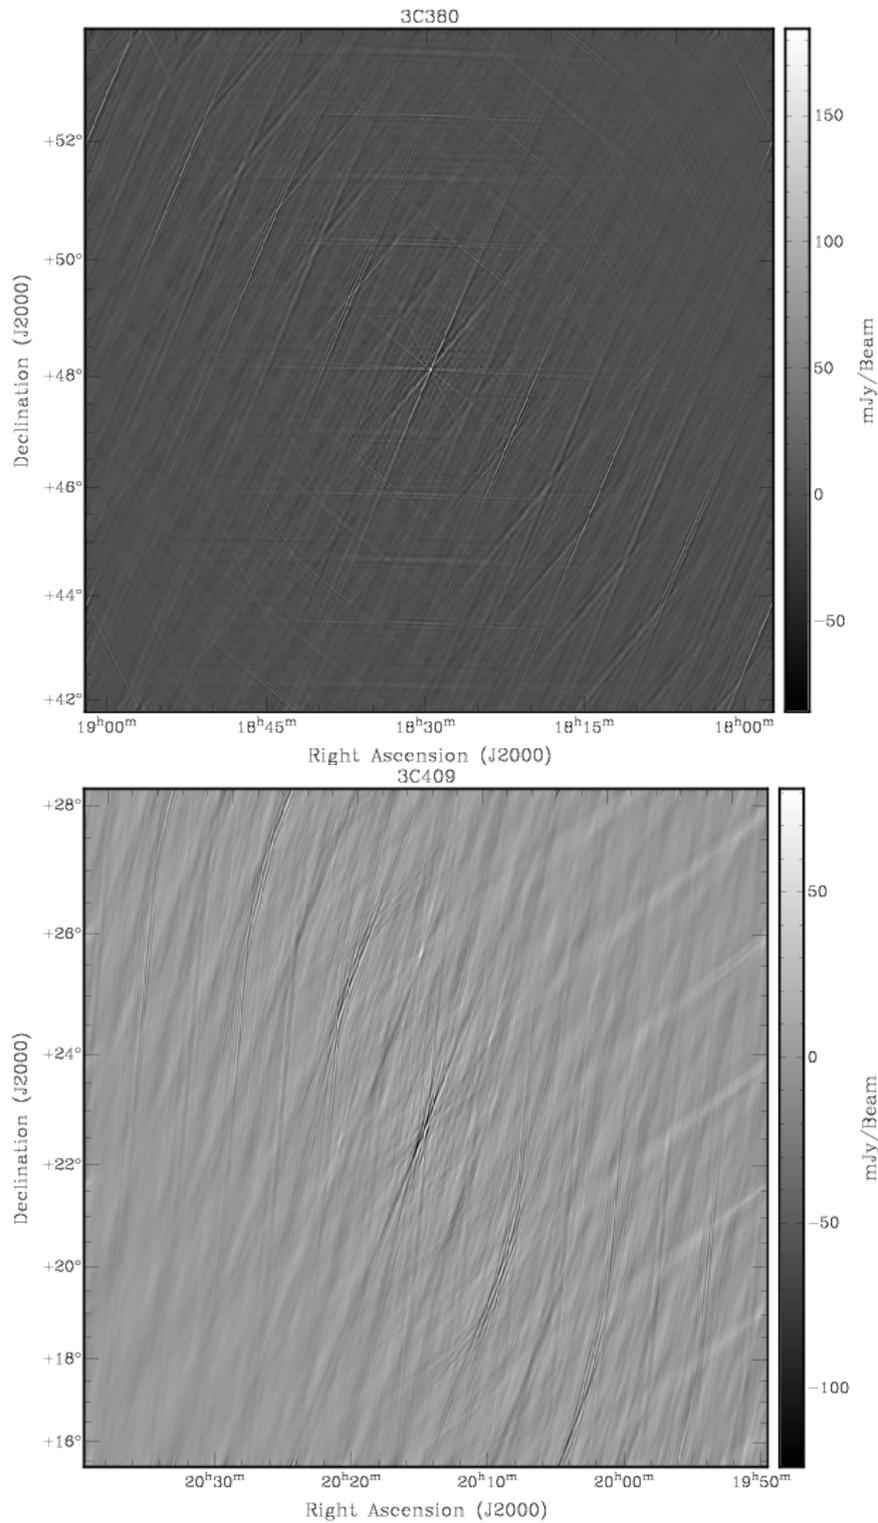
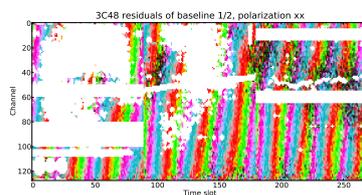
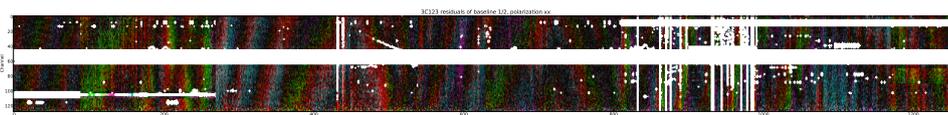


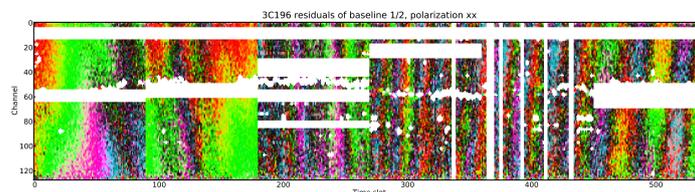
Figure A.6: *Dirty images of the residuals (corrected for modeled instrumental effects) of six fields with different greyscales (negative, flux is shown in black). The images have the same properties as Fig. 4.11. Dirty images of 3C147 are shown in Fig. 4.11*



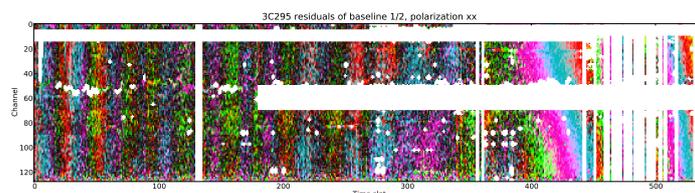
(a) 3C48



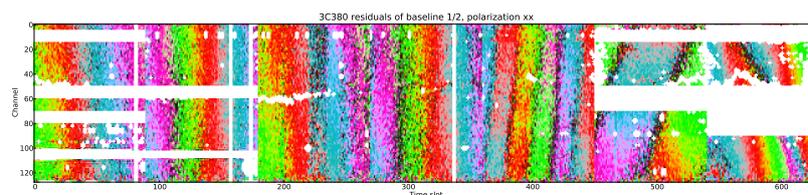
(b) 3C123



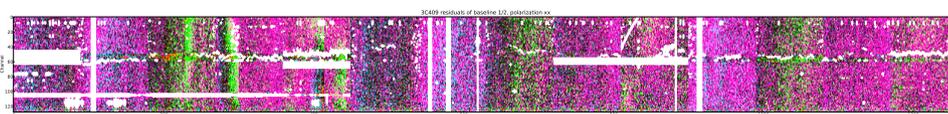
(c) 3C196



(d) 3C295



(e) 3C380



(f) 3C409

Figure A.7: Residuals for six fields on the same color scale for the fourth subband, baseline 1/2 and polarization xx . Residuals for 3C147 are shown in Fig .4.13. The color palette of appendix A.2 is used for displaying complex values.

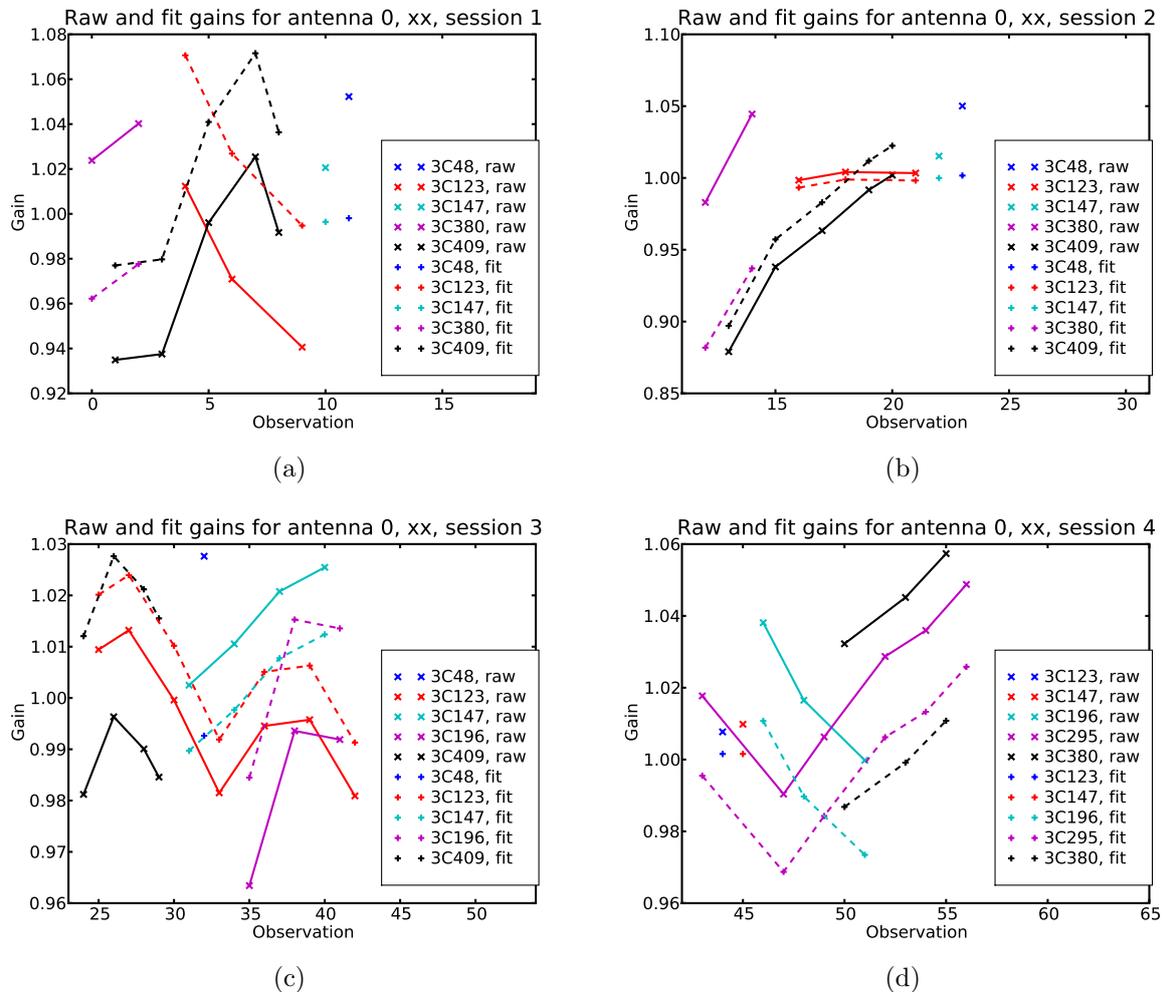


Figure A.8: Typical gain solutions solutions for the 12 observations of the seven sources during all four sessions. The diagonal crosses show the raw gain solutions and the upright crosses show the gain solutions after fitting the solutions of each field to one another. Gain solutions are connected with lines for fields with multiple solutions. The corrected gain solutions of all fields together look more like a smooth curve than the uncorrected ones. Gain solutions for the third session are also shown in Fig. 4.10

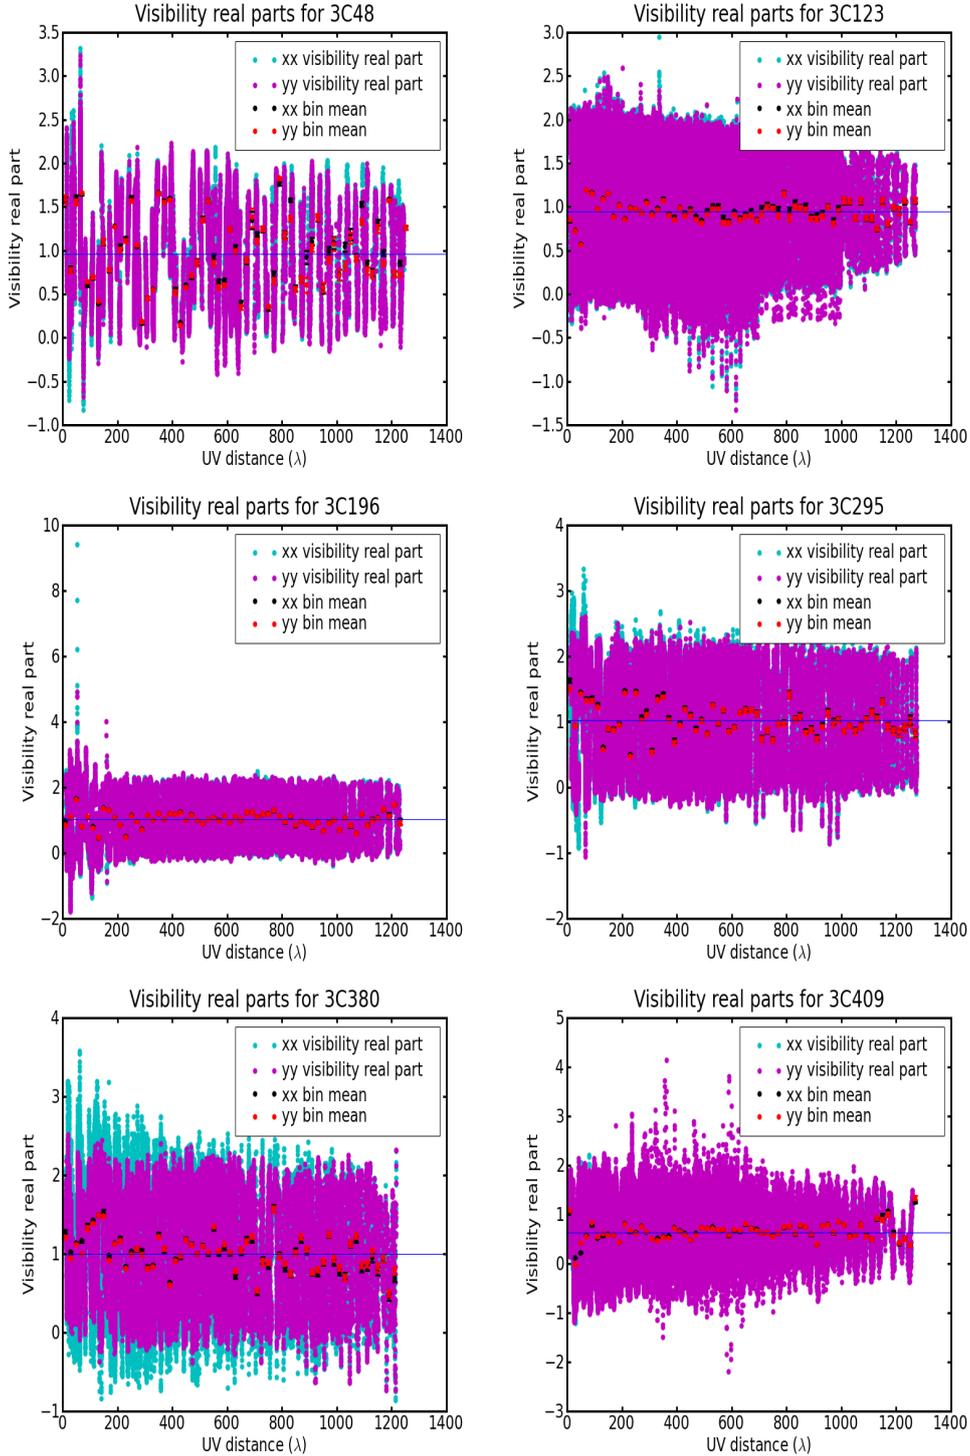


Figure A.9: Real parts of the 'corrected visibilities' (modeled off-axis fluxes subtracted and corrected for modeled instrumental effects) as a function of uv distance. Real parts were obtained by taking the mean of the real parts of channels 8-120. Cyan is the xx and magenta is the yy polarization. The black and red errorbars are means over 20λ bins for the xx and yy polarization, respectively. The blue horizontal line shows the radial weighted mean of the real parts. The real parts of the 'corrected visibilities' for other fields are shown in Fig. 5.1

Bibliography

- J. W. M. Baars, R. Genzel, I. I. K. Pauliny-Toth, and A. Witzel. The absolute spectrum of CAS A - an accurate flux density scale and a set of secondary calibrators. *A&A*, 61: 99–106, October 1977.
- A. S. Cohen, W. M. Lane, W. D. Cotton, N. E. Kassim, T. J. W. Lazio, R. A. Perley, J. J. Condon, and W. C. Erickson. The VLA Low-Frequency Sky Survey. *AJ*, 134:1245–1262, September 2007.
- S. L. Freeland. Solarsoft, 2009. <http://www.lmsal.com/solarsoft/>.
- E. W. Greisen. Aips memo 27: Non-linear coordinate systems in aips, 1983. <ftp://ftp.aoc.nrao.edu/pub/software/aips/TEXT/PUBL/AIPSMEMO27.PS>.
- J. P. Hamaker. Understanding radio polarimetry. IV. The full-coherency analogue of scalar self-calibration: Self-alignment, dynamic range and polarimetric fidelity. *A&AS*, 143: 515–534, May 2000.
- J. P. Hamaker. Understanding radio polarimetry. V. Making matrix self-calibration work: processing of a simulated observation. *A&A*, 456:395–404, September 2006.
- J. P. Hamaker and J. D. Bregman. Understanding radio polarimetry. III. Interpreting the IAU/IEEE definitions of the Stokes parameters. *A&AS*, 117:161–165, May 1996.
- J. P. Hamaker, J. D. Bregman, and R. J. Sault. Understanding radio polarimetry. I. Mathematical foundations. *A&AS*, 117:137–147, May 1996.
- F. J. Harris. On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform. In *Proc. IEEE, Volume 66, p. 51-83*, volume 66 of *Institute of Electrical and Electronics Engineers, Inc. Conference*, pages 51–83, 1978.
- S. M. Hill, V. J. Pizzo, C. C. Balch, D. A. Biesecker, P. Bornmann, E. Hildner, L. D. Lewis, R. N. Grubb, M. P. Husler, K. Prendergast, J. Vickroy, S. Greer, T. Defoor, D. C. Wilkinson, R. Hooker, P. Mulligan, E. Chipman, H. Bysal, J. P. Douglas, R. Reynolds, J. M. Davis, K. S. Wallace, K. Russell, K. Freestone, D. Bagdigian, T. Page, S. Kerns, R. Hoffman, S. A. Cauffman, M. A. Davis, R. Studer, F. E. Berthiaume, T. T. Saha, G. D. Berthiume, H. Farthing, and F. Zimmermann. The NOAA Goes-12 Solar X-Ray

- Imager (SXI) 1. Instrument, Operations, and Data. *SolPhys*, 226:255–281, February 2005.
- K. I. Kellermann, I. I. K. Pauliny-Toth, and P. J. S. Williams. The Spectra of Radio Sources in the Revised 3c Catalogue. *ApJ*, 157:1–+, July 1969.
- J. E. Noordam and O. M. Smirnov. Meqtrees, 2009. <http://www.astron.nl/meqwiki>.
- T. J. Pearson and A. C. S. Readhead. Image Formation by Self-Calibration in Radio Astronomy. *ARAA*, 22:97–130, 1984.
- R. B. Rengelink, Y. Tang, A. G. de Bruyn, G. K. Miley, M. N. Bremer, H. J. A. Roettgering, and M. A. R. Bremer. The Westerbork Northern Sky Survey (WENSS), I. A 570 square degree Mini-Survey around the North Ecliptic Pole. *A&AS*, 124:259–280, August 1997.
- G. A. Renting and M. A. Brentjens. Wsrt flagger, 2006. http://www.astron.nl/~renting/Handbook_WSRT_flagger.pdf.
- H. J. A. Rottgering, R. Braun, P. D. Barthel, M. P. van Haarlem, G. K. Miley, R. Morganti, I. Snellen, H. Falcke, A. G. de Bruyn, R. B. Stappers, W. H. W. M. Boland, H. R. Butcher, E. J. de Geus, L. Koopmans, R. Fender, J. Kuijpers, R. T. Schilizzi, C. Vogt, R. A. M. J. Wijers, M. Wise, W. N. Brouw, J. P. Hamaker, J. E. Noordam, T. Oosterloo, L. Bahren, M. A. Brentjens, S. J. Wijnholds, J. D. Bregman, W. A. van Cappellen, A. W. Gunst, G. W. Kant, J. Reitsma, K. van der Schaaf, and C. M. de Vos. LOFAR - Opening up a new window on the Universe. *ArXiv Astrophysics e-prints*, October 2006.
- R. J. Sault, J. P. Hamaker, and J. D. Bregman. Understanding radio polarimetry. II. Instrumental calibration of an interferometer array. *A&AS*, 117:149–159, May 1996.
- A. R. Thompson, J. M. Moran, and G. W. Swenson, Jr. *Interferometry and Synthesis in Radio Astronomy, 2nd Edition*. April 2001.