

Using Self-calibration to Model and Image the Virgo A Radio Source With the Low Frequency Array (LOFAR)

Alexander Spacek

April 1, 2014

Abstract

Through the National Science Foundation's International Research Experience for Students, four students from Arizona State University spent the spring semester of 2013 attending lectures each week to learn about radio astronomy, interferometry, and the Low Frequency Array (LOFAR), a radio telescope based in the Netherlands. In the summer of 2013 they spent two weeks working at Ruhr University Bochum in Bochum, Germany, and then five weeks working at the University of Hamburg in Bergedorf, Germany. They then spent seven more weeks continuing their work at Arizona State University. Their work involved using recent data from the LOFAR telescope on the science targets of the Coma Cluster and the Virgo A radio source. The work was divided up between the students, and this paper concerns the work on Virgo A.

This work consists of calibration, flagging, accounting for A-team sources (the bright northern-hemisphere radio sources such as Virgo A, Casiopeia A, Cygnus A, and Taurus A), modeling and imaging Virgo A, and imaging the wide field around Virgo A. The main part of the project involves trying out different methods of calibration, flagging, and imaging to try and get the best possible model and image of Virgo A. By combining 6 spread-out sub-bands to increase the frequency coverage, flagging the 4 antennas farthest away from the array core and then gradually adding two of them back in one at a time, and optimally selecting various calibration and imaging parameters, we are able to successfully get a model and image of Virgo A with a high resolution of 5" by 9".

1 Introduction

Looking at emission from the Universe at different wavelengths allows us to get a glimpse of the wide range of processes that produce radiation. Different processes tend to emit radiation at different wavelengths, or at least they emit in different ways at different wavelengths. Looking at the ultraviolet emission of a galaxy will likely give a very different picture than looking at the radio emission of that same galaxy. Compare, for example, the ultraviolet image of M31 from Gil de Paz et al. 2007 [3]¹ and the radio images of M31 from Berkhuijsen et al. 2003 [1]. For this reason it is important to design and use telescopes at all frequencies: radio, microwave, infrared, visible, ultraviolet, X-ray, and gamma ray.

The Low Frequency Array (LOFAR) radio interferometer telescope attempts to follow this philosophy in the low-frequency, long-wavelength, low-energy radio end of the electromagnetic spectrum. Radio covers radiation with wavelengths longer than about 1 mm or frequency less than

¹Image can be found at http://ned.ipac.caltech.edu/img15/GALEX_Atlas/MESSIER_031:I:FUV:g2006.jpg

about 300 GHz (in this definition we are including the microwave regime as well). LOFAR was made to exploit the relatively long-wavelength part of this range, and it can effectively operate within the wavelength range of 1.2–30 m or the frequency range of 10–240 MHz. See the introduction in van Haarlem et al. 2013 [7] for a clear overview of where LOFAR fits within the history of radio astronomy and what its scientific motivations are. For this paper I will mention that LOFAR has the potential to investigate, among other things, high-redshift 21 cm emission from the epoch of reionization, nanosecond radio flashes from ultra-high-energy cosmic rays, high redshift radio sources based on deep sky surveys, surveys of pulsars and cosmic radio transients, the detection of exoplanets, and the study of cosmic magnetic fields. To help LOFAR reach these potentials, it is important to develop a system of effective data-reduction techniques in order to be able to convert the original noisy, artifact-ridden LOFAR data into a science-ready image. For this paper we worked on doing so with a LOFAR observation of Virgo A.

The outline of this paper is as follows: in §2 I discuss the basics of radio astronomy, outline the LOFAR telescope, and describe our observation of Virgo A. In §3 I describe the self-calibration cycle we use, how we deal with interfering A-team sources, and the various fields we looked at. In §4 I present the various parameters we use and resulting images we create while trying to produce the best image (and corresponding model) possible. In §5 I summarize our results and briefly point out potential further work on this topic.

2 Radio Interferometry, LOFAR, and Virgo A

2.1 Why Interferometry?

Diffraction makes it so that the longer the wavelength we want to look at, the larger the telescopes we need to build. This is summed up by the Rayleigh criterion, which gives the relation between angular resolution (roughly the minimum angle two point sources need to be apart in order to be resolved) and the diameter of the antenna: $\theta = \lambda / d$. If we want to look at long wavelengths, such as those LOFAR detects, with good resolution and with single-dish telescopes, they need to be massive. To get an angular resolution of 15" at LOFAR's minimum wavelength of 1.2 m, a single dish would have to have a diameter of 16.5 km. This is unreasonably large, considering the largest single-dish telescope in the world is the Arecibo Observatory radio telescope with a diameter of 0.305 km.² A way around this is using interferometry to utilize a collection of scattered individual antennas to simulate a single dish with the same effective size as the longest baseline between two of the antennas.

For example, the maximum baselines of LOFAR's core stations, remote stations, and international stations are 3.5 km, 121 km, and 1158 km, respectively (the station types will be explained in §2.3). At 121 km, the remote stations have well over the effective diameter of 16.5 km needed to get an angular resolution of 15" at a wavelength of 1.2 m as mentioned in the above example. In this way, interferometry can be used to get good resolution at very long wavelengths that could never be reached any other way.

2.2 Basics of Radio Interferometry

In interferometry, two separate receivers record radiation from a source (see Figure 1). The physical distance between the two receivers means that the light has to travel farther to one of the receivers

²See <http://www.naic.edu/public/the.telescope.htm>




Figure 1: Simplified schematic drawing of a two-element interferometer.
 (from *Synthesis Imaging in Radio Astronomy II* [5])

than to the other (unless it is coming from directly overhead). This introduces a time delay τ_G in the radiation received between the two receivers, which corresponds to a phase difference between the output voltages from the receivers. These voltages are then multiplied together and integrated over time, and the final output voltage is only dependent on the amplitude and phase produced from the incoming radiation (for a given frequency). In an array of receivers, each set of receiver pairs is known as a baseline. Data from a single baseline at a given time for a specific wavelength gives information for a single point in the u - v plane, which is a Fourier transform of the actual image (l - m) plane. Imaging the sky is simply taking the inverse Fourier transform of all the u - v plane measurements from all of the different baselines. Equation (1) shows how the actual intensity distribution on the sky $I_\nu(l, m)$ relates to the measured spatial coherence function $V_\nu(u, v)$ through an inverse Fourier transform over u and v [6].

$$I_\nu(l, m) = \int \int V_\nu(u, v) e^{2\pi i(ul+vm)} du dv \quad (1)$$

The main reason it is difficult to get a perfect image is that the baselines are inevitably going to only partially fill in the u - v plane, leading to a lack of complete information needed for the inverse Fourier transform. The point of calibration and cleaning is to try and make up for this lack of information. Since what the telescope records is a Fourier transform of the actual sky, scales become flipped between the two domains. For example, smaller baselines are sensitive to larger scales in the sky and larger baselines are sensitive to smaller scales in the sky.

2.3 LOFAR Configuration

LOFAR is fundamentally based on two types of antennas, the high band antenna (HBA) and low band antenna (LBA), specifically HBA tiles and LBA dipoles. They both have dual polarization, meaning they are sensitive to two orthogonal linear polarizations (such as X and Y, see §2.4.1). The

HBA runs at 110–240 MHz and the LBA runs at 30–80 MHz. The antennas are put into groups called stations. There are core stations (CS), remote stations (RS), and international stations. The core and remote stations are all located in the Netherlands while the international stations are located in the nearby countries of Germany, France, Sweden, and the UK. There are currently 24 core stations, 16 remote stations, and 8 international stations. In the core stations there are two groups of 48 LBA each and two groups of 24 HBA each (known as “ears”). In the remote stations there are two groups of 48 LBA each and one group of 48 HBA. In the international stations there is a group of 96 LBA and 96 HBA. The minimum baseline length is 68 m, the maximum CS baseline is 3.5 m, the maximum RS baseline is 121 m, and the maximum international baseline is 1158 m [7].

2.4 Imaging Virgo A

2.4.1 LOFAR Observation

Virgo A is the fourth brightest radio source in the northern hemisphere. It is about 16.4 Mpc away and contains the galaxy M87 and its associated large-scale radio emission. This is due to synchrotron radiation caused by the energy transported through the intra-cluster medium by the jet and buoyant bubbles coming from the galaxy, which hosts an active galactic nucleus (AGN) resulting from an accreting supermassive black hole [2]. It is an interesting science target because of the extended radio lobes around the galaxy, extending up to 80 kpc. Studying this galactic and intra-cluster feedback from the AGN can help us understand the general processes of AGN feedback and the specific history of the M87 AGN. Virgo A is part of what is called the “A-team” of bright radio sources, which also includes Cassiopeia A (Cas A), Cygnus A (Cyg A), Taurus A, and Hercules A.

We look at Virgo A using an observation with the LOFAR HBA done during the night of May 8 to May 9, 2013. Since it is an HBA observation it spans the frequency range of 110–240 MHz. The output of this observation consists of data divided into many sub-bands, each representing a small frequency range. Each of these is divided further into 4 channels, each a small frequency range within the sub-band. The particular sub-band we first use is labelled *201* and was chosen to be at a frequency where the HBA have good sensitivity. It has channel 1 starting at 151.49 MHz and a channel width of 48.83 kHz, making the entire sub-band bandwidth 195.31 kHz. It also contains the polarization correlations XX, YY, XY, and YX. The data set started out with 32 core stations (16 stations, each with 2 “ears” of HBA), 13 remote stations, and 8 international stations.

A main goal of our work is to create an image of Virgo A at a frequency and resolution that has never been done before. To do this we do an initial setup and inspection of the data and then use a trial and error type method to try and produce the best model and corresponding images using self-calibration, flagging, and imaging. Another goal of our work is to use the best model of Virgo A we can create and use it to subtract Virgo A out of a much larger area of sky (from the same LOFAR observation) and create a wide-field image of radio sources around Virgo A. This is important because Virgo A is so strong in radio frequencies that if it isn’t subtracted out very carefully it obscures all of the sources around it.

2.4.2 Previous Virgo A Imaging

Virgo A has never really been looked at this closely in the full HBA frequency regime (110–240 MHz) before. The closest observation was done by de Gasperin et al. 2012 [2], who looked at Virgo A in the LBA frequencies of 15–77 MHz and the HBA frequencies of 115–162 MHz. Their HBA

observation had a maximum resolution of $19'' \times 14''$. The highest resolution obtained at frequencies near these was by Owen et al. 2000, who got a resolution of $7''$ at 327 MHz with the Very Large Array (VLA). As mentioned, our main goals are to image Virgo A at frequencies it has not been observed at yet (162–240 MHz) and at a better resolution than the previous best HBA resolution ($19'' \times 14''$), hopefully even approaching the overall high resolution of $7''$.

2.4.3 Data Reduction for LOFAR

In order to produce the best images possible, we use 3 main components of the LOFAR data reduction and imaging “pipeline.” These are BlackBoard Self-calibration (BBS), New Default Pre-Processing Pipeline (NDPPP) flagging, and the CASA tool `CLEAN`. Using these tools, we try to manipulate the quality of the resulting images by changing various parameters for BBS and `CLEAN` and flagging different data in NDPPP. We begin with the best parameters Francesco de Gasperin, the post-doc we were working with, had come up with in his years of working with the LOFAR data reduction pipeline. We calibrate, flag, and create an image using these initial parameters, and then we vary several specific parameters and repeat the process to see if the resulting image is improved. The main issues we run into are specific baselines having bad data, `CLEAN`-ing method, and the bright radio A-team sources Cas A and Cyg A.

2.4.4 Pre-Processing

The data set we start out with is a measurement set (MS) containing all the information produced by the observation in a format to be used with all of the LOFAR data reduction software. The MS comes pre-flagged by the program NDPPP based on “RFI excision” and other known issues from the time of observation, especially those related specifically to the workings of the telescope [5].

3 Methods

3.1 Data Inspection

The first thing we do with our pre-processed observation data is look at what we have to begin with. We inspect the details of the observation and the antennas used using the command `msoverview` (see pg. 10 of *The LOFAR Imaging Cookbook* [5]). This gives us information such as precisely when the observation took place, the coordinates of the target, the number of channels, the channel frequencies, the polarization correlations, and the name of every antenna and its coordinates. This is all information that is essential to know when going through with the data manipulation. We also use the command `casaplotms` to view the MS in the interactive `PlotMS` environment in CASA. This allows us to view plots of parameters such as amplitude vs. time, amplitude vs. u - v distance, and u - v coverage (see §4.1). We can view all the data at once or look at individual baselines, allowing us to inspect the baselines for suspiciously high or out of place amplitudes.

3.2 Self-calibration Cycle

3.2.1 Calibration

BBS is used for the calibration and simulation of LOFAR data. The data we have from LOFAR contains information from the actual sky, which is what we want, distorted with information from the Earth’s ionosphere and the telescope itself. The ionosphere creates phase-shifts in the radio

radiation that passes through it because of variations in the number of electrons in different parts of the ionosphere. On large scales these phase-shifts get averaged out, but on short scales they can have very significant effects on the data. This means that the longer baselines are most sensitive to ionospheric distortions [7] (since they correspond to the smallest scales on the sky), and this is the main reason why it is helpful to flag the farthest out antennas (see §4.2.2). The telescope introduces distortions due to characteristics such as the beam size and shape, which can amplify the signal from sources away from the target that you don't want to look at, and clock drift, a significant problem because timing is so essential in interferometry.

The goal of the self-calibration cycle is to get the best possible estimation for what the actual sky looked like during the observation. Something called a *measurement equation* is used to quantify the atmospheric and instrumental distortions. An example of a measurement equation is:

$$\vec{V}_{ij} = J_{ij} \vec{V}_{ij}^{IDEAL} \quad (2)$$

where \vec{V}_{ij} is the observed measurement for baseline ij , \vec{V}_{ij}^{IDEAL} is the ideal measurement, and J_{ij} represents all corruptions on baseline ij . J_{ij} can be factored into each separate corruption. These include polarization-independent multiplicative effects introduced by the troposphere, parallactic angle, effects introduced by properties of the optical components of the telescope, instrumental polarization response, electronic gain response, bandpass response, and baseline-based correlator errors [4]. From this we get a model of what the target is actually supposed to look like. Calibration is the process of finding out what parameter values need to be combined with the observation to get it as close to the model as possible [5]. Once data is calibrated we know what the specific atmospheric and instrumental distortions are and we can divide them out of our data to get as good as possible of an image of the actual sky. We use BBS as the first step of our self-calibration cycle to calibrate the data.

3.2.2 Flagging

NDPPP is used for the flagging and averaging of LOFAR data. It has the ability to do many tasks, including flagging, averaging, phase shifting, combining sub-bands, A-team demixing and subtraction, and filtering out baselines or channels [5]. Our main uses for NDPPP are both flagging data and combining sub-bands before starting the self-calibration, and flagging data within the self-calibration cycles.

As discussed in the previous section, we have to flag the farthest antennas (which logically contribute to the longest baselines) because they contribute the most to distortions in the data due to the ionosphere and timing issues for data traveling over longer distances. We first permanently flag all of the international stations because they are responsible for baselines much longer than any of the core or remote stations, and the distortions they introduce are much larger and harder to deal with than the scope of our project. Later we also experiment with flagging different combinations of the 4 farthest (from the LOFAR core) remote stations, sacrificing resolution for less distortions in our images.

We use NDPPP to combine sub-bands for various imaging tests we are performing. In order to get a better signal-to-noise ratio by increasing the amount of data we have, we combine 3 adjacent sub-bands, 4 channels each, into one MS file with one channel. In order to get better $u-v$ coverage by increasing our frequency range, we combine 6 sub-bands spread out within the full bandwidth into one MS file with one channel for each sub-band.

NDPPP is also used as the second step of our self-calibration cycle. After calibration the corrected data may still contain abnormally high amplitudes that were not accounted for, and we

try to eliminate these by flagging amplitudes over a certain threshold. This produces neater data to send to the next step in the self-calibration cycle, imaging.

3.2.3 Imaging

Once the data is calibrated and flagged, it is ready to be imaged. To do this, we use the CASA routine `CLEAN`. `CLEAN` creates an image and then deconvolves all of the sources out of it until only noise is left, resulting in an accurate picture of what all the sources looked like. In a single-scale run of `CLEAN`, the algorithm uses delta functions to model each source as a point source or collection of point sources. It calculates what each point source would have looked like with the telescope’s point spread function (PSF) and then multiplies it by a specified gain factor. It then divides out this source from the image, and the idea is to repeat this process until all of the sources in the image have been accounted for and only noise remains. This results in a representation of what all the sources looked like. In a multi-scale run of `CLEAN`, both delta functions and circular Gaussian functions are used in the deconvolution [4]. This helps model the extended, non-point source emissions much better since extended sources are more accurately modeled with extended Gaussian functions than with point-like delta functions.

During the calibration, the measurements in the u - v plane can be weighted in different ways. We use “briggs” weighting, which can be flexibly set to interpolate weights between “natural” weighting and “uniform” weighting. “Natural” weighting just weights each measurement by the noise in the data. This gives the highest signal-to-noise ratio since each measurement is used to its full potential. This also gives the worst angular resolution because shorter baselines are much more numerous than longer baselines and the longest baselines are sensitive to the smallest scales. “Uniform” weighting weights gridded cells in the u - v plane equally, which means longer baselines have equal influence with shorter baselines. This produces the best angular resolution since it emphasizes the longest baselines, but it produces poor signal-to-noise because it amplifies the gridded cells with the least amount of data in them. In briggs weighting, a robust parameter can be defined that runs between -2 and 2, where -2 is close to uniform weighting and 2 is close to natural weighting.

Our imaging is done in two steps. First we run 300 cycles of single-scale `CLEAN` in order to work out the very brightest points in the image. Then we run 10,000 cycles of multi-scale `CLEAN` to deconvolve the rest of the image. The results of a run of `CLEAN` are several files containing the telescope’s PSF, the residual that was left over after the deconvolution, the model, and the image itself. There is also the option of adding a mask to the image so that it only deconvolves sources within the mask. This is useful because it allows us to only deconvolve what we know are actual sources and not artifacts or distortions. We use CASA to create a mask around Virgo A and use it in the `CLEAN` self-calibration step for the 6 sub-band tests. Imaging using `CLEAN` was the third and final step in our self-calibration cycle. When doing multiple cycles of self-calibration, the output model will be used in the next cycle’s calibration so that it can build off of this cycle and ultimately create a better model and image itself. Then another cycle can improve upon that, and so on. If everything is working well, the models and images should start converging to the best possible quality we can get with our data and parameters. Then we can change the data and parameters, use the best model produced in our previous run, and start the self-calibration cycle again.

3.3 Clipping the A-team

A potentially major component of error in any radio observation done in the northern hemisphere is the group of A-team radio sources. To analyze and deal with problems in our data due to the

A-team, we used several pre-made scripts from the LOFAR imaging pipeline [5]. To broadly inspect potentially problematic sources, we use `plot_Ateam_elevation.py`. This creates a plot of elevation vs. time for the strongest radio sources (the A-team, Jupiter, and the Sun) during the time of our observation and tells us how far the sources are away from the target (see Figure 2 and §4.2.3). Once we have an idea of which sources will potentially be the biggest problems, we use `simulate.py` to simulate what data these sources would have produced in our observation. Then we use `compare.py` to tell us what fraction of our data would be affected by these sources. If enough data is messed up by these strong radio sources, it could be beneficial to try and remove the trace of these sources from the data. We use the program `Ateamclipper.py` to do this by flagging data that would be affected by model A-team source amplitudes greater than a specified threshold. We can then look at how much was flagged and try to make an image from the data to see if and how the image has improved by using the A-team clipping method.




Figure 2: Elevations of the strongest radio sources near Virgo A. (generated using the LOFAR program `plot_Ateam_elevation.py`)

4 Results

4.1 Initial Data

The international stations introduce complications that are very difficult to account for due to their extremely long baselines, such as timing issues and completely non-correlated atmospheric conditions. Our first step is to remove these international stations from our data. We do this using NDPPP by simply inputting only CS and RS and outputting what we input. Then we inspect the data by using the CASA command `plotMS` to plot things like amplitude vs. time (Figure 3, left), amplitude vs. u - v distance (Figure 3, right), and u vs. v (Figure 4). There are some drastic spikes throughout the data for certain antennae, especially towards the end of the observation. This can be seen in both core stations and remote stations (Figure 5).

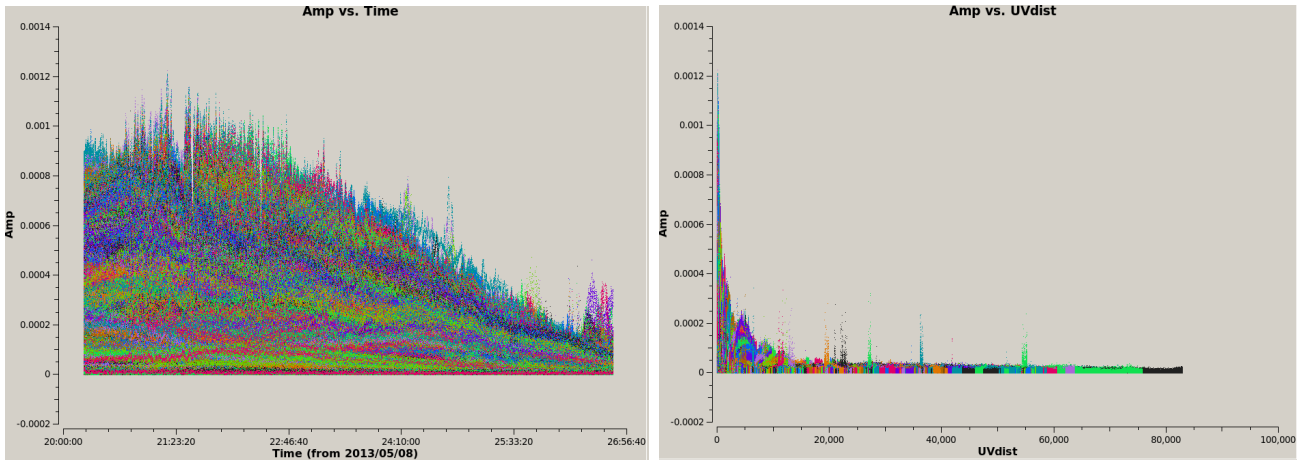


Figure 3: *Left*: Amplitude vs. time of all data, colored by baseline. *Right*: Amplitude vs. u - v distance of all data, colored by baseline. Both panels were generated using casaplotms.




Figure 4: u vs. v for all data, colored by baseline.
(generated using casaplotms)

4.2 Single Sub-band

Calibrating involves several steps. First, we take a starting model image provided by Francesco deGasparin (the post-doc we worked with) based on previous Virgo A work of his and fill in the “model column” of our MS. Francesco’s model is based on LBA data and it has a poor resolution of 20” which we hope to improve upon. We create a program to run multiple iterations of self-calibration. We have to set up a parameter set (parset) file to run BBS with.³ To start, we try 3 different methods of calibrating. One is with a u - v range of 250 wavelengths and above, one is with

³The complete set of parameters can be seen here:
<http://www.lofar.org/operations/doku.php?id=engineering:software:tools:bbsconfigurationsyntax>

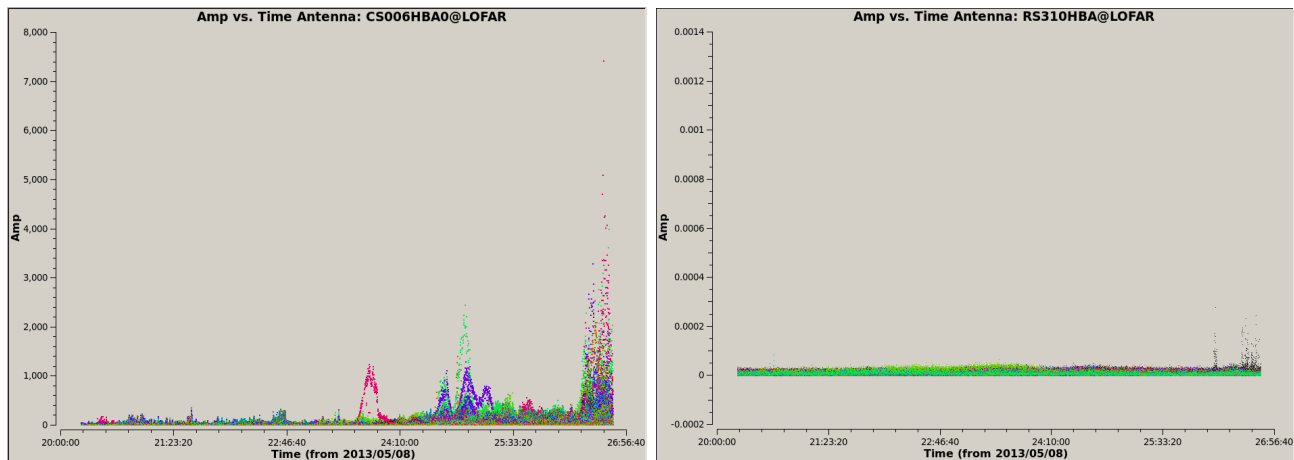


Figure 5: Amplitude vs. time. *Left:* Antenna CS006. *Right:* Antenna RS310.
Both panels were generated using casaplotms.

a $u-v$ range from 0 to 30,000 wavelengths, and one is without solving for the common rotation angle. The reason the $u-v$ range of >250 is used is to cut out the ear-to-ear baselines in the core stations, which cause problems in the data because of their extremely short distances from one another (as little as 68 m [7]). The reason for trying to remove the longest $u-v$ distances is to see if it helps remove the longest and most problematic baselines. The reason for solving without the common rotation angle is to save computation time. Using a common rotation angle involves figuring out how each HBA station is oriented with respect to the source and taking that into account when measuring things like polarization. The data is then cleaned. First, 300 iterations without multi-scale are done. Then, 10,000 iterations with multi-scale are done. To start out, we use “briggs” weighting with a robust parameter of 0.0.

4.2.1 Initial Calibration Methods

We start out just by running one self-cal cycle. This produces one image each for the >250 wavelengths and 0-30k wavelengths $u-v$ ranges, and the no rotation angle setup (Figure 6). Looking at these, the 250 image looks the best and the 30k and no rotation angle images look okay but with more artifacts. We decide to use >250 wavelengths in the calibration from now on.

4.2.2 Initial Imaging Methods

Using the >250 $u-v$ calibration, we try removing different stations to see how the images improves. We try cleaning with antenna RS310 removed and the image improves noticeably (Figure 7, top left). We do the same with RS509 and get the same result (Figure 7, top right). Cleaning with both RS310 and RS509 removed gives the best result so far (Figure 7, bottom). Next we decide to start the whole self-cal cycle of calibration, flagging, and cleaning with the previous 2 antennas removed. We run this for 7 cycles and then used the resulting model (produced from cleaning) to start a new set of self-cal iterations. This was in order to replace the poor quality original model we were using that was based on a completely different data set. We then ran it through 50 cycles of self-calibration (Figure 8).




Figure 6: Virgo A after 1 cycle of self-calibration with 1 sub-band and a robust parameter of 0.0.

Top-Left: Using a u - v range of >250 wavelengths. *Top-Right:* Using a u - v range of 0 to 30,000 wavelengths. *Bottom:* Using a u - v range of >250 wavelengths without solving for the common rotation angle. All panels were generated using casaviewer

4.2.3 Dealing With the A-team

By plotting the elevations of the strongest radio sources, including the A-team, the Sun, and Jupiter, as well as their angular distances from our source, we can see that Cas A and Cyg A both significantly increase in their elevations towards the end of our observation (recall Figure 2). Therefore, our next step is to simulate the effects of Cas A and Cyg A and see if we can subtract them out of our data and maybe account for the poor data near the end of our observation.

To simulate what effects Cyg A and Cas A would have had on our observation, we run an available LOFAR program called `simulate.py` which simulates a mock measurement set for the source we specify, during the observation time we specify, at the right-ascension and declination we specify. We get output MS files for both Cyg A and Cas A, and using CASA plotMS again we can see what they would have looked like during our observation (Figure 9). We can see that both sources increase in amplitude towards the end of the observing time, especially Cyg A.

Next we run another LOFAR program called `compare.py` to see how much the simulated data would have affected our actual data. According to that, Cyg A affected 6.60% of the XX polarization and 8.47% of the YY polarization, and Cas A affected 3.57% XX and 6.02% YY. These could be fairly significant effects. We decide to use the LOFAR program called `Ateamclipper.py` to try and flag Cyg A and Cas A from our data. First we fill the model column of our MS file with the




Figure 7: Virgo A after 1 cycle of self-calibration with 1 sub-band and a robust parameter of 0.0 using a u - v range of >250 wavelengths. *Top-Left*: With antenna RS310 removed. *Top-Right*: With antenna RS509 removed. *Bottom*: With antennas RS310 and RS509 removed.

All panels were generated using casaviewer

appropriate model of Cyg A or Cas A using BBS calibration, and then we use `Ateamclipper.py` to flag our data based on the models. For Cas A about 9% of our data is flagged, and for Cyg A about 10% is flagged. These numbers are fairly high, so it might be possible that we could significantly improve our data by flagging these two A-team sources. We decide that we will first try to calibrate Virgo A without dealing with these A-team sources to create a base image to compare any poten-




Figure 8: Virgo A after 50 cycles of self-calibration with 1 sub-band and a robust parameter of 0.0, using a $u-v$ range of >250 wavelengths with antennas RS310 and RS509 removed.
(generated using casaviewer)




Figure 9: Amplitude vs. time after running `simulate.py`. *Left:* For Cas A. *Right:* For Cyg A. Both panels were generated using `casaplotms`

-tially improved images with, and then cut out the A-team sources to see if it improves the image.

We also try clipping the A-team sources of Cas A and Cyg A on the >250 $u-v$ calibration with the two antenna removed (Figure 10). It doesn't make any noticeable improvements in the image, so we decide to not worry about clipping the A-team from now on.

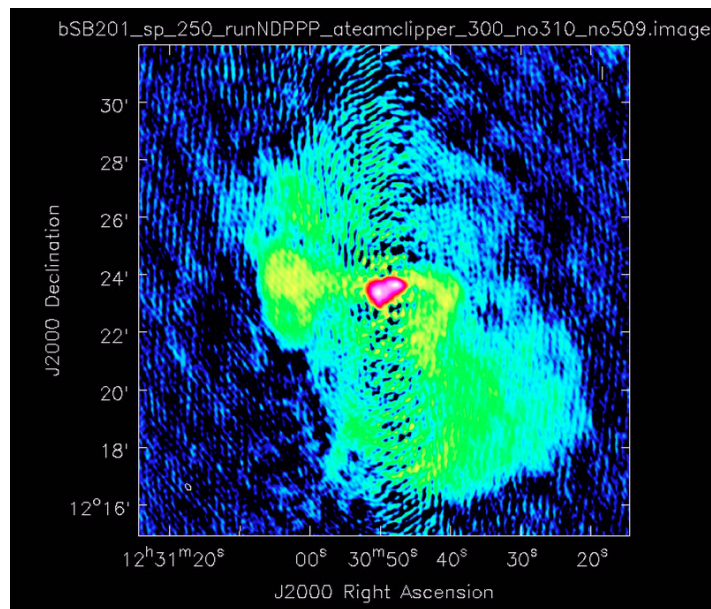


Figure 10: Virgo A after 1 cycle of self-calibration with 1 sub-band and a robust parameter of 0.0, using a $u-v$ range of >250 wavelengths with antennas RS310 and RS509 removed and Ateamclipper applied for Cas A and Cyg A. (generated using `casaviewer`)

4.3 Three Adjacent Sub-bands

We decide to increase the size of our data set by combining 3 adjacent sub-bands into one MS file. The reason for this is to increase our signal-to-noise ratio by having more data, though we have to be careful because we are increasing the bandwidth we are using, and calibration solutions are frequency dependent. We combine three sub-band MS files into one channel in one MS file. We run

one set of self-cal cycles with all antennas, and one set with RS509 and RS310 removed, like before. We run both sets through 50 cycles of self-calibration (Figure 11).

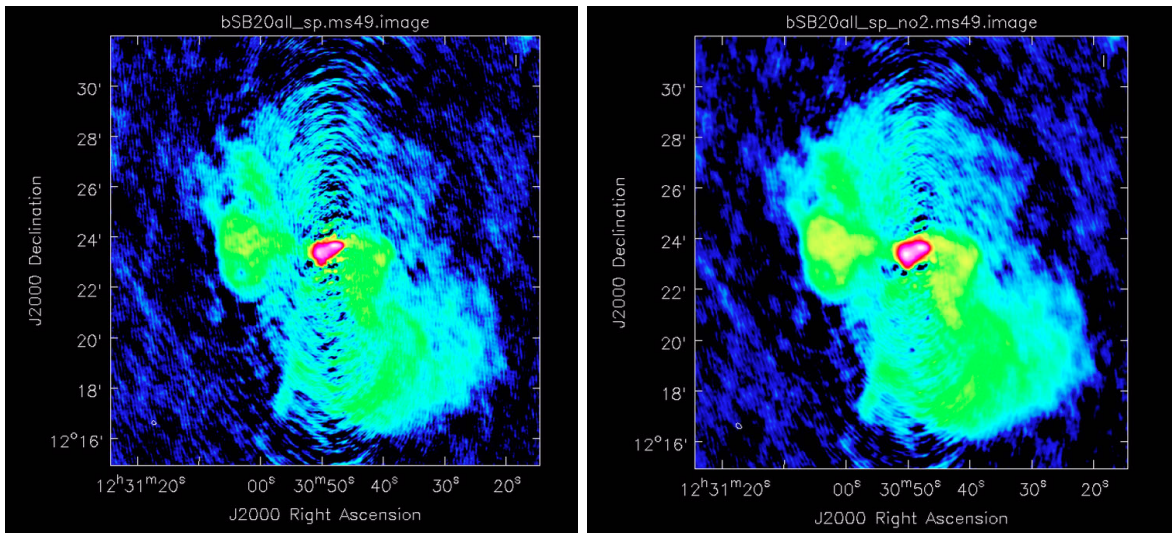


Figure 11: Virgo A after 50 cycles of self-calibration with 3 adjacent sub-bands and a robust parameter of 0.0. *Left:* With all antennas. *Right:* With antennas RS310 and RS509 removed. Both panels were generated using casaviewer.

4.4 Six Interspersed Sub-bands

We further decide to increase the wavelength range of our data set by using 6 sub-bands each spaced 10 MHz apart. The u - v coordinates of a baseline are wavelength dependent since they are, by definition, measured in terms of wavelength [6]. Therefore using a greater variety of wavelengths will increase the u - v coverage of the data, resulting in a better image. We use sub-bands at 122.1 MHz, 131.8 MHz, 141.6 MHz, 151.4 MHz, 161.1 MHz, and 170.9 MHz. We combine these 6 sub-bands into one MS file with 6 channels, each representing a sub-band.

4.4.1 Removing Four Antennas

We then remove the 4 farthest away remote stations, RS208, RS310, RS508, and RS509, in order to get the best image possible at the sacrifice of resolution. We plan on going through cycles of self-cal until the cleaned images converge to a certain quality, and then we can use that model on a new cycle of self-cal that has one less station removed. We run this first setup through 5 cycles of self-calibration (Figure 12).

4.4.2 Removing Three Antennas With Varied Robust Parameters

Next, we try removing only the 3 longest antenna, RS310, RS508, and RS509. We use CASA to draw a mask around Virgo A to be used in cleaning from now on, which we expect to improve the results. We decide to try using different robust parameters to see which will work best, so we use 5 different sets of self-calibration cycles with robust parameters of -1.0, -0.5, 0.0, 0.5, and 1.0. We run the 1.0 robust through 15 cycles of self-calibration, and it does not produce a recognizable image of Virgo A (Figure 13, bottom left). This means something probably went wrong, and we would have to change the calibration and cleaning parameters in order to get an image. We run the 0.5 robust




Figure 12: Virgo A after 5 cycles of self-calibration with 6 spread out sub-bands and a robust parameter of 0.0, using a u - v range of >250 wavelengths with antennas RS208, RS310, RS508, and RS509 removed. (generated using casaviewer)

through 4 cycles of self-calibration and it produces a recognizable image, but the image quality is horrible (Figure 13, bottom right). Having a positive robust parameter does not appear to work well. We run robust -1.0 through 20 cycles (Figure 13, middle left), robust -0.5 through 42 cycles (Figure 13, middle right), and robust 0.0 through 20 cycles (Figure 13, top). Robust -1.0 produces a pretty poor image. Robust -0.5 produces a decent image. Robust 0.0 produces the best image, and we decide that it will be best to continue using a robust parameter of 0.0.

The angular resolution of an image in casaviewer is given by an ellipse in the bottom left of the image. By measuring that ellipse and confirming it by looking in the CASA log file for that particular run of CLEAN, we can see what angular resolution we were able to achieve in the image (see left image of Figure 14). For the run with a robust parameter of 0.0 (and 3 antennas removed), we get down to a resolution of $6''$ by $11''$.

4.4.3 Removing Two Antennas

Next, we try removing only the 2 longest antennas, RS310 and RS509, and using a robust parameter of 0.0 (Figure 15). We run this for 13 cycles of self-calibration. Using the same method for finding the angular resolution as mentioned in §4.4.2, we find we have achieved an angular resolution of $5''$ by $9''$ (see right image of Figure 14).

To get an idea of what we're looking at, compare our result shown in Figure 15 with a previous image of Virgo A shown in Figure 16 (top) with the structures labeled, as well as an image with the physical scale shown in Figure 16 (bottom). We get a good view in our image of the various major structures. The ~ 5 kpc wide inner cocoon can be seen in the center with a dark hole just below it that is most likely an artifact from the cleaning procedure, which has a hard time handling the extremely bright object. This inner cocoon contains a jet pointed towards the northwest (north = up, west = right). There are two large, ~ 40 kpc haloes, one to the north and one to the south of




Figure 13: Virgo A with 6 spread out sub-bands using a u - v range of >250 wavelengths with antennas RS310, RS508, and RS509 removed. *Top*: After 20 cycles of self-calibration with a robust parameter of 0.0. *Middle-left*: After 20 cycles of self-calibration with a robust parameter of -1.0. *Middle-right*: After 42 cycles of self-calibration with a robust parameter of -0.5. *Bottom-left*: After 15 cycles of self-calibration with a robust parameter of 1.0. *Bottom-right*: After 4 cycles of self-calibration with a robust parameter of 0.5. All panels were generated using casaviewer.




Figure 14: 6 sub-band, robust = 0.0 zooms to measure the resolution for *Left*: 3 antennas removed and *Right*: 2 antennas removed. Both panels generated using casaviewer.

the center. They are connected by two large flows. The east flow is oriented nearly east-west and ends in a pair of bright lobes. The west flow quickly twists as soon as it leaves the central region and heads to the south, where it spreads out into filaments. While the east flow seems contained within its lobes at the end, the west flow appears to fill out the halo with its plasma filaments [2]. With our resolution we can start to see a little bit of structure in the filaments to the north and to the south, as well as emerging structure in the west flow. We can also clearly see that the flows are all confined quite effectively within the outer halo boundary.

4.5 Widefield Imaging

Another line of imaging we decide to do is widefield imaging around Virgo A, using data from the same observation we have been using previously. Like we have done with Virgo A, we take 3 sub-bands and combine them into one MS file with one channel. We also remove the same 4 farthest-away antennas. We calibrate using BBS, and then use CASA to subtract Virgo A from the image using our best model of Virgo A produced from all the other runs. We then clean it and inspect the image, which is pretty good besides a fairly prominent cross-hatch pattern across the whole image. We try clipping the A-team the same way we've done with Virgo A previously, and then cleaning again (Figure 17). The cross-hatch pattern is completely removed. It appears that for the widefield image the A-team plays a prominent role in image quality.

5 Summary, Conclusions, and Further Work

We try many different methods for calibrating, flagging, and imaging Virgo A. Our initial tests show that the best u - v range is one that cuts out the smallest baselines and uses the rest. Then, we discover that the longest baselines are contributing the greatest distortion to our data in the imaging. We have already removed all of the international stations, and we find that removing the two longest remote stations greatly improves the results. We also run the Ateamclipper program

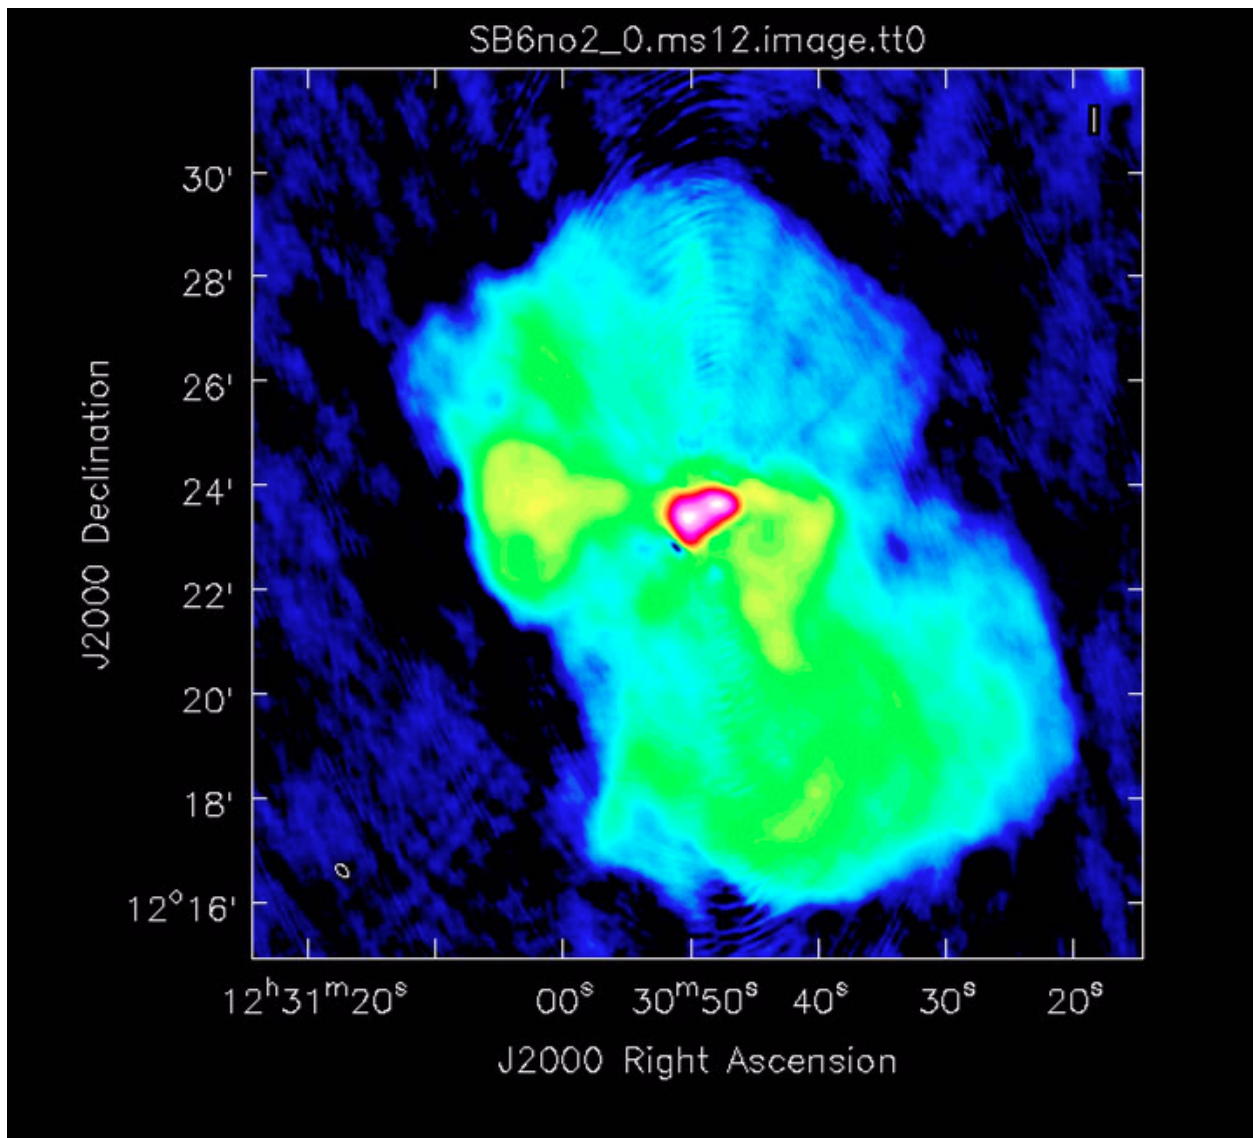


Figure 15: Virgo A after 13 cycles of self-calibration with 6 spread out sub-bands and a robust parameter of 0.0, using a $u-v$ range of >250 wavelengths with antennas RS310 and RS509 removed. (generated using casaviewer)

at this time and find that it doesn't noticeably clean up the image. We then attempt to increase our signal strength by using more data spread across multiple nearby sub-bands, but the images have not improved much. We next try to increase our wavelength coverage and the corresponding $u-v$ coverage by using 6 sub-bands spread out across the full bandwidth, and this provides the best results. Going from there, we first remove the four farthest antennas and run self-calibration cycles until the images appear to be converging in quality. Then we use the best model from that to run self-calibration cycles on a MS with only three antennas removed. At this point we try out several different robust parameters in the Briggs weighting scheme in CLEAN, and find the robust parameter of 0.0 that we have always been using to be the best. We then go to the next step and use our best model to run self-calibration cycles on a MS with only the two farthest remote stations removed. With this we have achieved the best image (and corresponding model) with a resolution of $5''$ by $9''$.





Figure 16: *Top*: Virgo A with its various structures labeled. *Bottom*: Virgo A with its angular and physical scale shown. Taken from De Gasperin et al. 2012 [2]




Figure 17: Widefield with Virgo A subtracted out after 1 cycle of self-calibration and a robust parameter of 0.0, using a u - v range of >250 wavelengths with antennas RS208, RS310, RS508, and RS509 removed. *Left:* Without Ateamclipper applied. *Right:* With Ateamclipper applied for CasA and CygA. Both panels generated using casaviewer.

As mentioned in §2.3.2, before this work the highest LOFAR HBA image resolution of Virgo A was $14''$ by $19''$. Our resolution has greatly improved this value. Also, the highest ever resolution of Virgo A near these frequencies was $7''$, and we have nearly matched this with our ellipse of $5''$ by $9''$. If our project was taken further and more antennas added back in to the data, we could conceivably surpass this value. As mentioned in §4.4.3, we are starting to get good images of the extended structures in Virgo A, including the central source, flows, haloes, and the outer boundary. With further improvements in resolution we could get an even better picture of the finer details within the structures.

We also use one of the models of Virgo A we have created to subtract Virgo A out of the whole observation field and make a widefield image of the radio sources around Virgo A. To do this we follow our previous Virgo A results by removing four antennas, and the resulting image still has prominent cross-hatch artifacts. We decide to run the Ateamclipper program on the widefield image, and in this case it actually does effectively improve the image by removing those artifacts.

Our results can play an important role in guiding future work on imaging Virgo A and other LOFAR targets. From our work it appears that Virgo A is bright enough and far enough from the other A-team sources in the sky that it can be imaged well (at least initially) without having to worry about subtracting out the other A-team sources. If a large field around Virgo A is being used, though, the A-team probably still have a significant effect on the data. When using LOFAR to view Virgo A, it will likely be necessary to remove some of the farthest remote stations and slowly build them back in to the image as the models are improved through cycles of self-calibration.

There is potential for further work, which could continue gradually improving our model and image of Virgo A. Ultimately we could have all antennas included to get the best possible resolution and have a reliable and accurate model of Virgo A for future work. There is also the potential to make a giant mosaic of flanking wide fields around Virgo A which were all observed at the same time as the Virgo A data we've been using. This undertaking would need the best possible model of Virgo A so that it could be subtracted out of all the flanking fields. It would include the problem

of subtracting Virgo A from an image where it is not in the center using a model where it is in the center. There are also fairly bright artifacts around the edges of some of the flanking fields due to bright radio sources located just outside of the full images that would have to be accounted for and removed.

References

- [1] Berkhuijsen, E. M., Beck, R., & Hoernes, P. 2003, *Astronomy and Astrophysics*, 398, 937
- [2] de Gasperin, F., Orru, E., Murgia, M., et al. 2012, *ArXiv e-prints*, arXiv:1210.1346
- [3] Gil de Paz, A., Boissier, S., Madore, B. F., et al. 2007, *The Astrophysical Journal Supplement Series*, 173, 185
- [4] Ott, J., & Kern, J., eds. 2013, *CASA User Reference & Cookbook*, 4th edn. (National Radio Astronomy Observatory)
- [5] Pizzo, R. F., ed. 2012, *The LOFAR Imaging Cookbook: Manual Data Reduction With the Imaging Pipeline*, 12th edn. (ASTRON)
- [6] Taylor, G. B., Carilli, C. L., & Perley, R. A. 1999, *ASP Conference Series*, 180
- [7] van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, *ArXiv e-prints*, arXiv:1305.3550