NUI Galway
OÉ Gaillimh

Observing with I-LOFAR

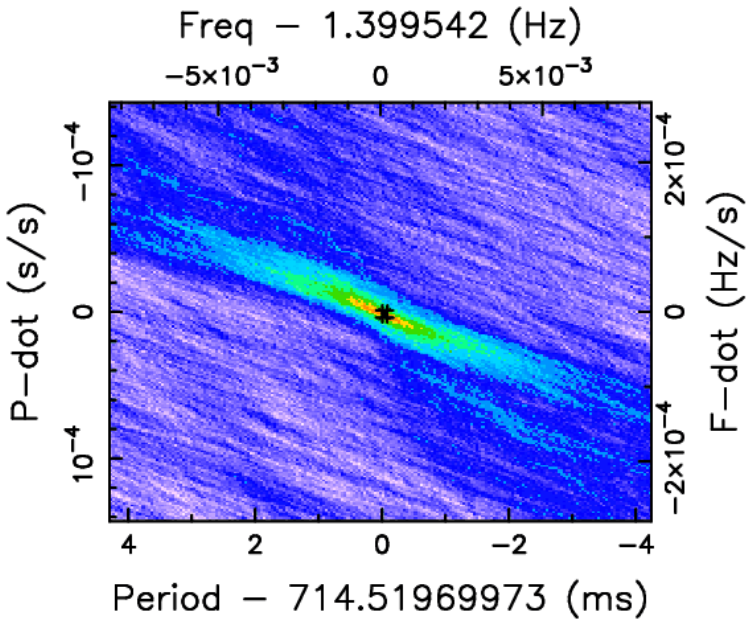

by

Owen Higgins

4BPT2 – 4th Year Physics

Supervised by: Dr Matt Redman

April 2019

Abstract

The I-LOFAR radio interferometry telescope is a research-grade telescope situated right here in Ireland and available to researchers from NUI Galway. It has been recently extended with a new high time-resolution processing cluster for imaging transient phenomena, REALTA. This project examines the capabilities of both the I-LOFAR station and the REALTA cluster including the feasibility of their different possible observation techniques and modes; with a view towards a student being able to record data from the telescope and analyse it. All-sky radio maps are produced from the full antenna array and it is shown that this kind of imaging is well-developed and reasonably practical for students to perform. Measurement of pulsar radio emission is attempted also using REALTA for recording and processing. It is shown that this type of observation is already possible and areas for future development are identified, with the observations taken in this report contributing to the understanding of the processing methods. Guidance is given on foot of these results for areas of future observation including consistent monitoring of existing pulsars. In addition, an understanding of the science behind and the hardware and software of the station, necessary for planning and making sensible observations with this unusual telescope, is developed and presented here.

Table of Contents

1 Introduction 4

1.1 Objectives 4

2 Background theory 5

2.1 Radio astronomy 5

2.2 Interferometry 7

3 Instrumentation and methods 10

3.1 I-LOFAR 10

3.1.1 Hardware and software 10

3.1.2 Capabilities 12

3.2 REALTA 13

3.2.1 Hardware and software 13

3.2.2 Capabilities 14

3.3 All-sky imaging 14

3.4 Transient imaging 16

4 Results and discussion 18

4.1 All-sky images 18

4.2 Pulsar plots 20

5 Conclusions and future work 23

6 Bibliography and references 24

6.1 Books used 24

6.2 Reference list 24

7 Acknowledgements 25

A. Appendix: Pulsar altitude and flux plotting 26

Appendix: Plagiarism statement

1 Introduction

This project, undertaken between January and March 2019, aimed to employ the new radio interferometry telescope I-LOFAR for observation and data analysis. I-LOFAR – the Irish Low Frequency Array – is a newly-constructed telescope situated in County Offaly in the centre of Ireland, and is accessible to researchers from a number of universities and institutes across Ireland, including NUI Galway. On top of this a new supercomputing cluster, REALTA, has recently been completed and added to the telescope by NUI Galway, Trinity College Dublin and University College Cork. This system was completed only in 2018 and its potential uses for data processing and analysis are just beginning to be developed. In this project, some of the capabilities of the telescope and this new high-performance cluster were examined through a number of observing runs done in early March 2019 and further processing and investigation of the resulting data.

I-LOFAR is a research-grade telescope in the form of a radio interferometer, allowing Irish astronomers to examine the radio universe from the point of view of north-western Europe. It also forms part of the wider International LOFAR Telescope network, being one of several stations spread across Europe. With the addition of REALTA – the REALtime Transient Acquisition Cluster – the telescope is uniquely capable of high time-resolution processing required to analyse objects such as pulsars and fast radio bursts.

This report will detail some background to radio astronomy, interferometry and the telescope itself, before outlining the types of imaging that we undertook for this project, and then our results and conclusions from these. This information is based upon sources listed in the bibliography and references, as well as upon conversations with researchers involved in I-LOFAR and REALTA, to whom we are indebted.

1.1 Objectives

The primary objective of this project was simply to examine the capabilities of the telescope, determining what types and modes of observation were feasible to perform with it. This was done by trying to obtain and analyse data products; a particular focus was on REALTA and transient imaging. Secondary objectives included examining how practical the station was to use, especially for students and in projects such as this one.

2 Background theory

A large portion of the work on this project was undertaken to build an understanding of the science behind the telescope, and especially of the particulars of its hardware and software. This part of the work was necessary to develop observation plans and proceed with data acquisition. We include an overview of this information as an introduction to the later discussion of methods and results.

2.1 Radio astronomy

While astronomy is the oldest of sciences non-optical observations are relatively new. Radio astronomy was first performed by Karl Jansky at Bell Labs in the 1930s, by chance: Jansky was investigating sources of radio interference, and picked up on a repeating signal which he worked out was from an interstellar radio source in the Milky Way passing in front of his antenna every night. While Jansky did no further work in the field, some research was performed during World War II, and from the 1950s onwards academic researchers began to direct their efforts towards detecting more emission sources and mapping the radio sky.

Radio astronomy is performed in much the same way as optical astronomy, by pointing a telescope at the sky and receiving light from celestial sources. The much longer wavelength of radio light, however, necessitates differences in collectors and detectors. The typical radio telescope is an antenna, comprised of a large parabolic dish reflecting waves to a small receiver above it. As with optical telescopes, the angular resolution of the image formed by the telescope depends on the aperture size *D*, which for a radio telescope is the diameter of the dish, according to the Rayleigh criterion:

$$\theta = \frac{1.22\lambda}{D} \quad (1)$$

A telescope measuring at 1.42 GHz (the wavelength of the Hydrogen hyperfine transition), to obtain a resolution of just 1 arcminute, would require a dish 880 metres in diameter. While no telescopes exist with dishes this large, several do have dishes in the range of hundreds of metres.

Much radio astronomy is done at wavelengths from 3 metres to 30 cm and below – more commonly expressed as frequencies of 100 MHz to 1 GHz and above. These wavelengths are best collected with dishes. Low-frequency radio astronomy, from 10 MHz

up to 100 MHz, can be performed with smaller directional wire antennas much like the television antennas visible on homes in Ireland. As I-LOFAR is sensitive in this low range, it exclusively uses antennas of this type, and doesn't require large dishes.

The usefulness of radio astronomy is in its ability to image phenomena invisible to the naked eye. There are many regimes of astronomy beyond optical; radio is notable for being able to image the active nuclei of galaxies, remnants of supernova explosions, the cosmic microwave background radiation, and compact objects like – famously – pulsars. Familiar objects such as the Sun can have broad-band emission that includes radio waves. I-LOFAR is particularly suited for surveying these low-frequency emissions on an on-going basis.

Radio astronomy is hampered by the prevalence of radio-frequency interference (RFI), mostly due to the extensive use of radio waves for communications technology. This interference sometimes appears in sky maps, and it must be mitigated (by cleaning the obtained images) or avoided (by observing in different frequency sub-bands). During scouting for the construction of the new LOFAR telescope priority it was discovered that Birr was quite low in RFI and so quite suitable for the placement of the telescope [I-LOFAR 2017]. A similar problem arises from the Earth's ionosphere blocking large bands of radio emission. Radio wavelengths above about 30 metres must be imaged by satellite telescopes above the atmosphere, but everything below this down to the far infrared is transmitted.

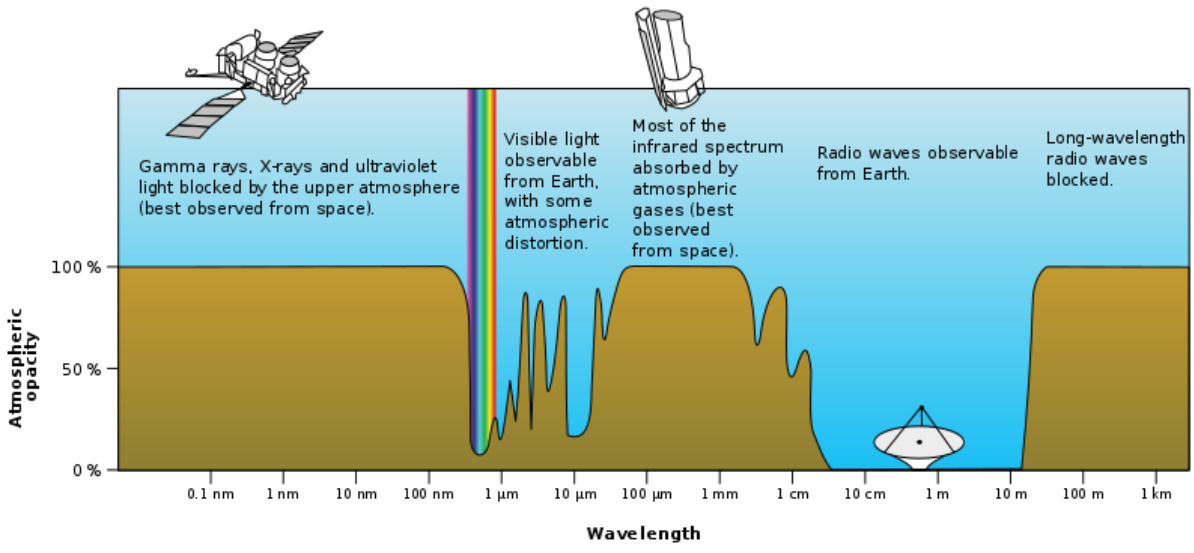

Figure 2.1: Opacity of the Earth's atmosphere across the EM spectrum

2.2 Interferometry

The difficulty of constructing antenna dishes large enough to obtain satisfactory resolution turned astronomers to alternative technologies. Interferometry was first applied to radio astronomy by Ryle and Vonberg with their two-element array in 1946, and was developed further with calibration and image-processing techniques, tracking antennas and Earth-rotation synthesis [Thompson pp. 13-14]. Some interferometers such as the VLA in the USA use a few large dish antennas for interferometry; LOFAR stations use a large number of much smaller antennas, exchanging individual resolution for more baseline pairs.

The basis of interferometry as employed by LOFAR stations is the baseline pair: an image formed by two individual antennas separated by some distance. Each antenna receives radio emission from the same source in the sky, recording it as a complex-valued voltage which is characterised by its phase and amplitude as a wave. *Interference* between these two waves cancels out differences and preserves the matching signal, resulting in a wave that represents the combination of the two signals, and should therefore accurately represent the single radio source from which the signal was received.

The important parameter in the case of interferometry like this is not the size of the individual elements, but the straight-line distance between them, called the *baseline*. Often the baseline is measured in wavelengths, given by the actual distance divided by the receiving wavelength. This baseline then determines the telescope's angular resolution. The size of the elements does however have an impact on how faint of objects can be imaged.

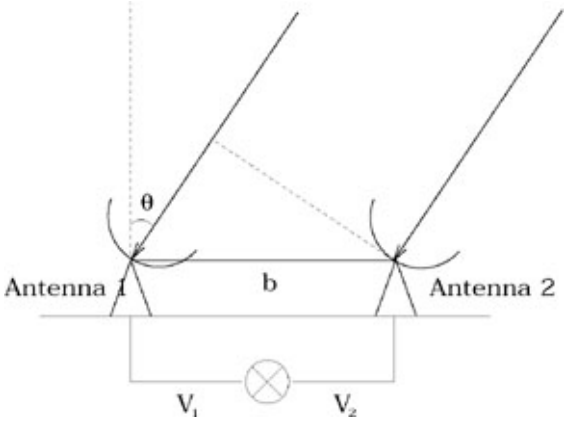

Figure 2.2: Two-element interferometer [Birkinshaw and Lancaster 2004]

The source is at an angle ϑ from the zenith. V_1 and V_2 are the complex signals.

The overall signal resulting from the interference of the two elements' signals is a complex-valued number (accounting for phase and amplitude of the signal). It is called a complex visibility, or cross-correlation[1]. The coordinates of this visibility are *u* and *v* on a flat plane called the *u, v*-plane, where

$$u = \frac{a}{\lambda} \quad v = \frac{b}{\lambda} \quad (2)$$

with *a* and *b* the coordinates of the actual antenna locations in real space, divided by wavelength λ. As a result *u* and *v* are in units of wavelengths.

It can be seen as in chapter 5.2 of Burke *et al*. [2010] that these correlations are the Fourier transform of the brightness distribution on the sky, of the imaged source. That is, $V(u, v) \Leftrightarrow B(x, y)$. As a result, we can recover the sky image by simply performing a Fourier transform on the received signal; Figure 2.3 demonstrates this for a two-element interferometer. The shape of this sky image is known as the 'beam'.

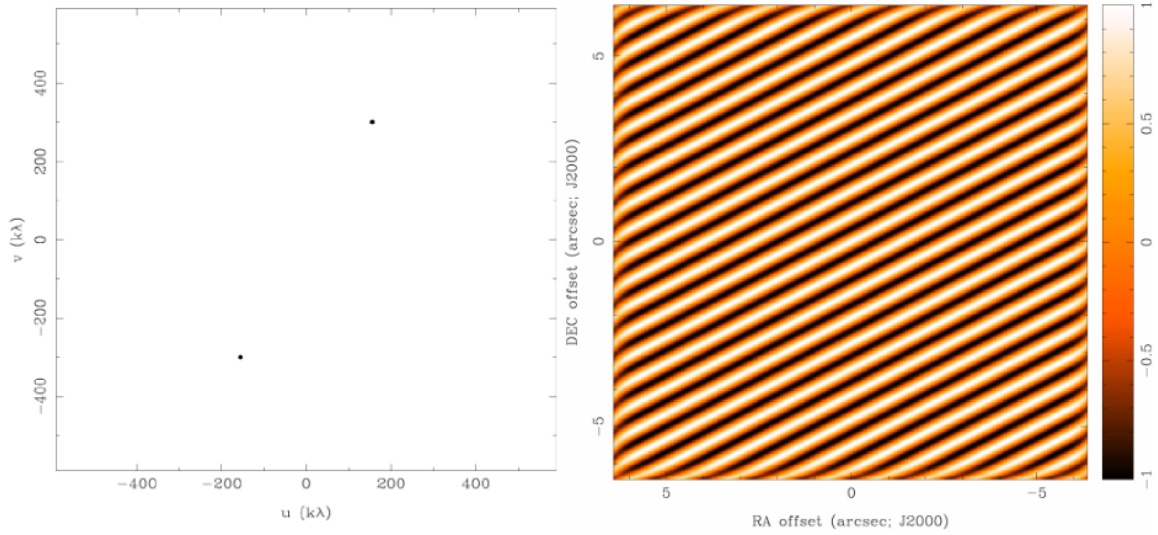

Figure 2.3: Fourier transform of visibilities in (u, v) to brightness in (x, y) on the sky [Isella 2011]

A single observation by this two-element interferometer produces only a single value in the *u*, *v*-plane, and a Fourier transform of one point is clearly not very useful. In order to reproduce the ideal PSF of a single source in the Fourier plane, the entire *u*, *v*-plane should be covered. We can approximate this very well by taking a large number of measurements and combining them using the technique of aperture synthesis.

[1] Strictly speaking cross-correlations are between voltage signals in the time domain, and complex visibilities between signals in the frequency domain, but both analyses are equivalent [Burke 2010 p. 77].

I-LOFAR is a good example of an *aperture-synthesis array*. Using its 96 antennas (either HBA or LBA), sources are observed by many two-element interferometers, formed from every possible pair of antennas. The number of pairs available is given by

$$\frac{N^2 - N}{2} \quad (3)$$

for an *N*-element array [Burke 2010 p.97], where the –*N* term comes from pairing antennas with themselves, resulting in so-called 'auto-correlations' that give brightness but no location information. I-LOFAR can have at maximum 4,560 correlations if all antennas in an array are operational. The coverage of the *u*, *v*-plane can be further increased by simply waiting as the Earth rotates, moving the antennas under the source and allowing them to take extra measurements from slightly different places. The *u*, *v*-plane is filled out by large ellipses as Figure 2.4 shows.

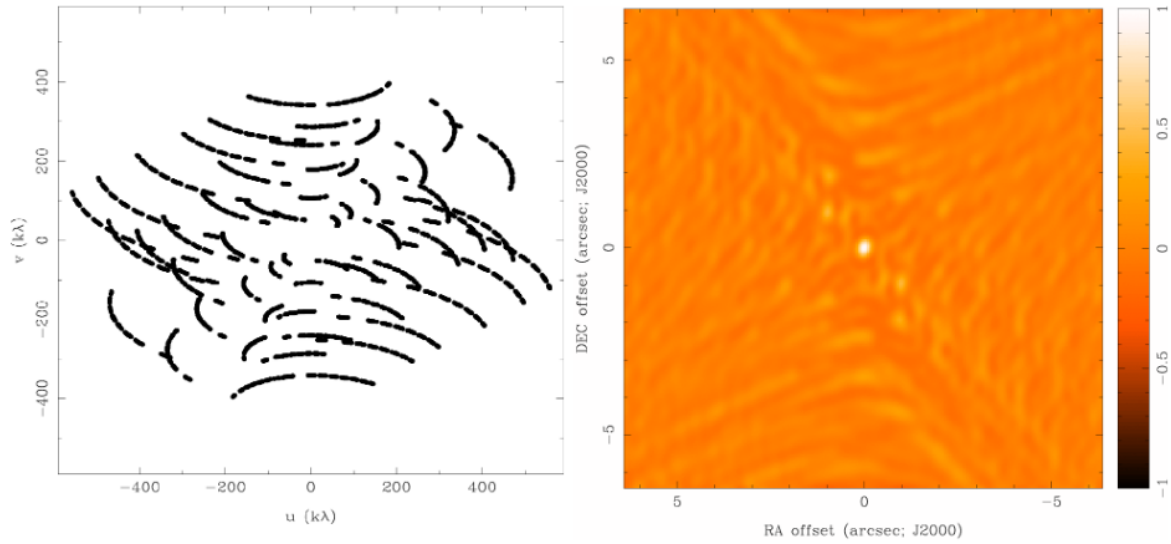

Figure 2.4: antennas in the u, v-plane drawing out arcs as the Earth moves, and a better Fourier image resulting from the increased coverage [Isella 2011]

The difference in signals between the two antennas of the two-element interferometer is caused by the time difference in signals reaching each. This delay can be eliminated, or it can be varied to modify the shape of the beam. In particular an artificial electronic delay can be introduced to cause constructive interference in the signals received from a particular point on the sky, and destructive interference for locations outside of that. This shapes the beam to focus on a particular target, and is known as 'beamforming'. I-LOFAR uses beamforming heavily to image targets other than just the entire sky, as outlined in section 3.1.2.

3 Instrumentation and methods

The focus of this project is the I-LOFAR radio interferometry telescope, and the REALTA high time-frequency processing cluster attached to it. Understanding the operation of both the station and the REALTA processing was vital to proceeding to observation, as well as itself being a goal of the project, and so a degree of that understanding is reproduced here.

3.1 I-LOFAR

The Irish Low Frequency Array telescope consists of a standard setup of high- and low-band antennas, connected to a variety of machinery which converts complex analogue voltages to digital signals and integrates these over 1-second or longer periods while recording the data. It has the same setup as all ILT stations across Europe and when in international mode it transmits this data via fibre link to LOFAR's data processing centre at Groningen, NL. In local mode the data is stored on-site to be processed and analysed. The means by which the station is controlled, and data is recorded, are outlined below; the vast majority of this information, where not otherwise referenced, is per the LOFAR Station Data Cookbook by Virtanen [2018].

3.1.1 Hardware and software

I-LOFAR's setup as standard consists of three elements: the low-band antenna array (LBA antennas); the high-band antenna array (HBA tiles), and an RF-shielded container containing all the processing hardware and computers. The LBA antennas are dipole wire antennas standing on metal grid bases, sensitive from 10 – 90 MHz, while the HBA tiles are bow-tie antennas housed in plastic-covered tiles, sensitive from 110 – 270 MHz. (The intermediate range is used for FM radio broadcasting.) Both contain signal amplifiers and coaxial cable connections to the container, while the HBAs also have analogue beamformers to supplement the RCU beamforming.

The antennas are connected to receiver units (RCUs) in the container, one RCU for each antenna and 192 total. The RCU perform important functions including input selection, signal amplification and frequency band filtering on the input signal; they then sample the data with 12-bit A/D converters, at either 200 or 160 MHz clock speed (the latter only in

mode 6). The combination of frequency band and sampling frequency is selected by choosing a single 'RCU mode', detailed in Table 3.1.

Mode	Input	Frequency range (MHz)	Passband (MHz)	RCU filter band (MHz)
0	*Off*	–	–	–
1	LBA	10 – 90	10 – 80	0 – 100
2	LBA	30 – 90	30 – 80	0 – 100
3	LBA	10 – 90	10 – 80	0 – 100
4	LBA	30 – 90	30 – 80	0 – 100
5	HBA	110 – 190	115 – 185	200 – 100[2]
6	HBA	170 – 230	170 – 220	160 – 240
7	HBA	210 – 270	215 – 260	200 – 300

Table 3.1: Available receiver, or RCU, modes

Mode 0 turns off the RCUs, modes 1 through 4 use the LBA antennas, while mode 5 through 7 use the HBA tiles. The passband above is the actual band that is imaged.

The RCUs are in turn connected to 24 Remote Station Processing (RSP) boards. These have two basic functions. First, they perform the Fourier transform on the signals to return them to a sky image, as outlined in 2.2, and in doing so divide the signal into 512 narrow sub-bands. The sub-bands are narrow enough (~200 kHz) that they can be used for beamforming, so that is done here, with each desired 'beamlet' corresponding to one sub-band. Secondly the RSP boards calculate the actual data products, which are either sub-band statistics (SSTs), beamlet statistics (BSTs), or 'crosslet' statistics (XSTs). These are described below. After this, the data is ready for local analysis, or transportation to Groningen for the ILT via fibre link.

The final component of the station hardware is the Local Control Unit (LCU). All functions and capabilities of the station can be controlled from here as described in the following section, including remotely via a `ssh` connection to the machine.

[2] The filter band and therefore signal spectrum is inverted in this mode, but it is re-inverted in later processing, so this has no net effect.

3.1.2 Capabilities

The I-LOFAR station operates in two basic modes: international and local. In the former it is used as part of the ILT and all data is transmitted back to Groningen for processing; in the latter, its use is at the discretion of researchers from NUIG, UCC, TCD and others. We employed a few hours of the weekly local time allocation to make our observations. In all other senses the two modes are identical, and both have the same capabilities (except that the international telescope has much, much longer baselines).

The capabilities of the station can be reduced to the three types of data products that it outputs – SST, BST and XST – as well as additional hardware on stations, such as REALTA. Sub-band statistics are the most basic data product that the station can produce. They simply output an average power in each sub-band produced by the polyphase filter [Virtanen 2018], producing a plot of intensity against frequency for a chosen band. It is possible to plot this in one dimension (frequency), or in two dimensions of frequency and time, with intensity represented by colour. SSTs are useful for measuring the frequency response of the telescope, the presence of RFI, and taking observations of objects where an x-y map of the object is unimportant.

BSTs are similar in nature. As described at the end of 2.2, the telescope beam is beamformed by introducing artificial electronic delays to the RCUs. `beamctl` has arguments for specifying the coordinates for pointing the beam; The telescope will track these while the beam is active. The command must specify the same number of sub-bands as it does beamlets, but the sub-band numbers and beamlet numbers do not have to match, and the latter usually just count from 0 to 487 (the max. number). While normally all the beams formed are pointed in the same direction, to measure a single target across a wide frequency band, it is fully possible to specify multiple different beamlets or groups of beamlets simultaneously, and so image multiple objects at once; the beams can also be set to turn off after a delay, and be restarted to image a new object, while other beams continue. As Virtanen [2018] discusses, this is not an intuitive process, but it can be done. The beamlet statistics themselves are essentially the same as SSTs, giving average power in each sub-band, the difference being that the sub-band is only measuring a region of the sky and not all of it.

Crosslet statistics, also known as array covariances, calculate the cross-correlations between all antenna pairs in the array. The data consists of a complex sample from all antenna pairs, in one sub-band at a time. This gives an image of the whole array's received signal, for making maps of the radio sky. It can be computed quite quickly depending on the desired resolution.

3.2 REALTA

The REALtime Transient Acquisition Cluster is a supercomputer built by NUIG and UCC to perform high time-frequency processing of data from I-LOFAR. Based on a similar system built for the UK's LOFAR station at Chilbolton named ARTEMIS, REALTA consists of four machines built by UCC and one by NUIG [I-LOFAR 2018*a*]. Because REALTA is non-standard to LOFAR stations, it has an unusual processing path and its uses and software have not yet been developed fully. During this project we had the opportunity to get involved in these investigations and examine its current feasibility for observations.

3.2.1 Hardware and software

LOFAR stations, while not naturally supporting instrumentation such as this, are easily extensible with custom third-party hardware.

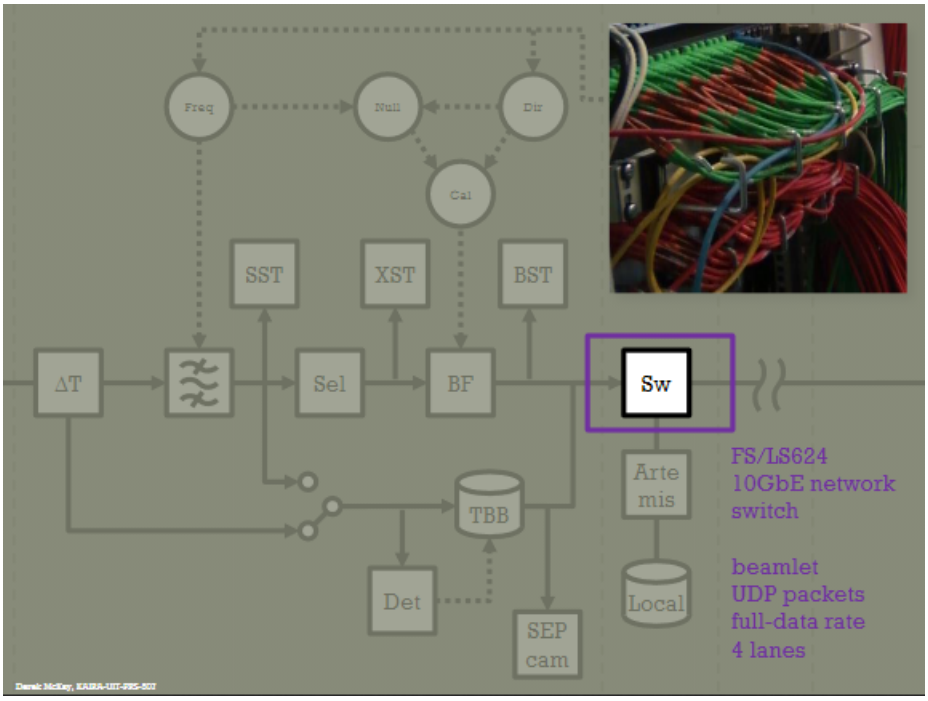

Figure 3.2: Schematic diagram of processing data path [McKay 2018]

A 10 Gb/s switch allows instruments to collect beamlet data from I-LOFAR at the full-data rate, that is, without the data being integrated or binned. This means the data can be sampled to a much higher time-resolution, essentially as high as desired, but it also means that vastly more data is produced per-recording. The data packets are output across 4 lanes, where lane 1 starts with the lowest quarter of sub-bands and lane 4 ends with the top quarter. REALTA is connected to this switch.

REALTA has no standard software either, but much like the LCU, shell scripts are executed on the system to tell it to record data. The recording must be synchronised with the data output from the LOFAR station, however, which means that scripts need to be executed on both machines at once, and the recording on REALTA must wait for the beams to start up on I-LOFAR. Currently this is done by running a wrapper script remotely that executes the scripts on both machines with appropriate delays.

3.2.2 Capabilities

In principle REALTA is capable of recording data at high time-resolution (millisecond and beyond). This is suitable for observing astronomical objects and phenomena that vary with high cadence or are very brief, collectively known as transients. However, a pipeline for data recording and processing must be developed for any desired observation types, mainly through writing scripts to run on REALTA. The pipeline for recording and processing pulsar data is an example of this and one which has come quite far, as will be described in sections 3.4 and 4.2 – I-LOFAR with REALTA is capable of imaging pulsars right now. It is certainly possible that similar pipelines could be developed for alternative transient phenomena such as fast radio bursts or SETI activities, and these could share data recording scripts and steps once these have been developed fully.

3.3 All-sky imaging

For this and the next section, we aim to outline how we proceeded with our observations, so that it can be replicated more easily by readers of this report and future users of the telescope.

While investigating possible observations that could be made with I-LOFAR, Peter Gallagher suggested to us that his project student David McKenna had completed a project on all-sky imaging using I-LOFAR, and that the scripts he used to process this data were

available online [McKenna 2018]. It was decided that this would serve as a good investigation of the telescope's capabilities in a standard, well-travelled area, although there was possibility for variation on our part too.

I-LOFAR takes all-sky images by recording XST data using the `rspctl --xcstatistics` command. A few parameters such as integration time are given as parameters while others are given to other commands. No pointing or beamforming is done; the field of view is from the zenith straight above the telescope down to ~15° from the horizon. The station can only record XSTs for one sub-band at a time; it must be made to loop over all, or a selection of, sub-bands. We observed every tenth sub-band for 2-second integrations, between sub-bands 10 and 80, for an hour; we did this in mode 3 and mode 5.

In general all-sky imaging is reasonably simple to do. The steps for a typical observation would be:

- Decide on observing frequency (hence receiver mode) and on sub-bands to observe in that mode (all, or a subset of them).
- Decide on observation duration and integration time per-image. (For example, ours was an hour long and two-second integrations.)
- Write an observing script that calls the correct `rspctl` commands for these parameters. It should also set the software level (to 2).
- On the LCU, ensure the station is idle (no other observations are running), then execute the script.
- Wait. The `rspctl` commands will output data to the LCU directly.
- When complete, the data should be transferred off of the LCU.

Regular observers would have a template observation script for this purpose into which they can fill the desired values for sub-band and RCU mode etc., making this process very simple after the first few observing runs.

After obtaining the all-sky data, it can be processed quite simply by means of a Fourier transform, and then plotted on the x-y (or, RA-DEC) plane. We used David McKenna's code for this, which uses a Discrete Fourier Transform-based algorithm to create circular maps out of the sky image. The module included with it some sample observational data, which we processed to learn the use of the software. We reproduce a plot of that data below to briefly explain the form of the output plots.

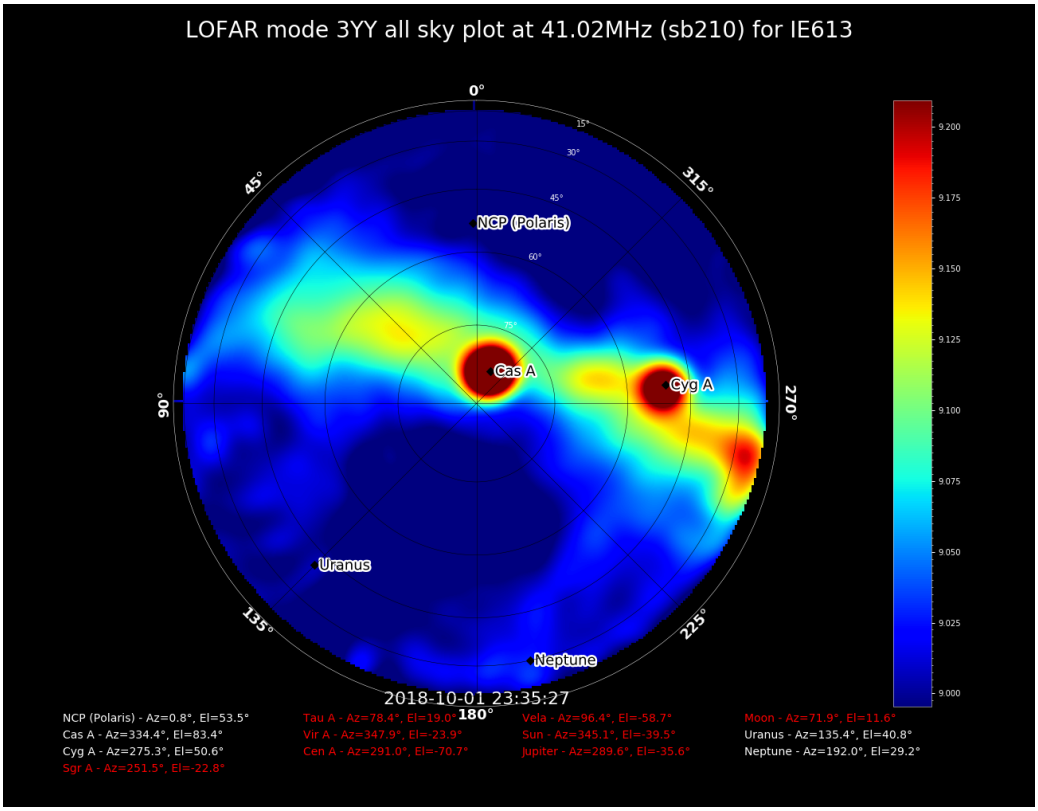

Figure 3.3: Output plot typical of our all-sky image code. Data per McKenna [2018]

The plot is annotated with, along the top: mode and polarisation, frequency and sub-band, and station ID; along the bottom, a list of known celestial objects including common strong radio sources; on the plot itself, the locations of any of those objects if visible, and the degrees of RA and DEC. Intensity of radio emission is represented by colour, where the colourbar along the right has units of mJy. Two strong radio sources are apparent in this image, Cassiopeia A and Cygnus A.

3.4 Transient imaging

From the outset of this project, we sought to determine whether it was currently feasible to use REALTA for transient imaging, especially as part of a short student project, and if so what could be imaged. On inquiry with researchers including Joe McCauley and Peter Gallagher we learned that work was ongoing to develop a pipeline for imaging of pulsars, comprising data recording and analysis; pulsar 'first light' on I-LOFAR was achieved in only September of 2018 [I-LOFAR 2018*b*]. We decided to attempt recording some pulsar data using REALTA and see if it could be processed to obtain coherent results.

Currently the pipeline for use with pulsar data (and how we used it) is as follows. REALTA as connected to the station switch receives UDP full-data packets from beamformed data; as a result, I-LOFAR must be commanded to record beamlet statistics (which are immediately discarded) for REALTA to obtain data. This is done with the `beamctl` command and we wrote observation scripts to do so; a script must be executed on I-LOFAR and on REALTA concurrently for the recording to succeed. Our imaging was done at a roughly 0.1-millisecond time-resolution, across the maximum number of 488 beamlets, divided across the four lanes as standard. All in all four data products were output, each ~11 GB in size.

Before analysis this data had to be processed by converting it to 'filterbank' format, a simple 2D stream of radio frequency channels versus time [Lorimer 2008]. Scripts for this purpose were provided by Joe McCauley, but the process is very processor-intensive, and only-suitable to be done on REALTA. We chose to only process lane 2 data, as lanes 1 and 4 have some empty sub-bands at the top and bottom respectively of their bands, and lane 3's band suffers from some known RFI.

A standard pulsar analysis library called PRESTO [Ransom 2015] is used to examine the data. Through our investigations of the capabilities of this library, and inquiries with other researchers at I-LOFAR, we learned there are two methods to analyse pulsar data with it. The period and dispersion measure (DM) can be given and the data folded immediately, or the data can be checked and cleaned of RFI and searched for pulse candidates before folding it with no given parameters. The latter method currently does not work well on LOFAR data but the former has shown results and we used it on our data.

The steps taken to measure and analyse pulsar data with I-LOFAR currently look something like this:

- Decide on a pulsar to image, accounting for its brightness in the low-frequency domain and its visibility at the chosen observing time. How we selected targets to image is outlined in appendix A.
- Write an observing script to perform the beamforming on the station and record BST statistics. It should set the correct software level (3) first and also send pointings to the RCUs and, if in use, to the HBA tiles.
- Write a separate script to perform the recording on REALTA concurrently with the recording on the station. Account should be made for the time needed to start up beams and other processes.

- On the LCU, ensure the station is idle (no other observations are running), then execute the script.
- Wait while the data is recorded to REALTA. Monitor the logs for packet loss (the number of packets is important to know later).
- Once the data is recorded, process by filterbanking each individual lane, specifying the number of packets to process, the packet format, and the top sub-band of the lane. (An alternative script may take different parameters here or calculate them itself.)
- Optionally, use PRESTO to search for RFI and eliminate it from the filterbank files.
- Optionally, use PRESTO to search for pulse candidates.
- Use PRESTO to fold the data and search for pulses and period.

It is probably not feasible to automatically filterbank the data immediately upon observing it due to the high processing requirements on REALTA, without significant improvements in efficiency of the scripts used to do so.

4 Results and discussion

At the conclusion of the project, we had obtained a large amount of all-sky data and a number of pulsar plots. We present some of these plots here as illustration of the results we obtained from these observations, along with some discussion of what we can conclude from them.

4.1 All-sky images

As mentioned in 3.3, we obtained 61 sky images in each sub-band and two observation modes, for a total of 854 images. Fortunately, there is not significant difference between the individual images in each mode, so we can present a small number of images from each to illustrate the results we obtained. The mode 5 observations were taken first, and the mode 3 observations several hours later, but on the same evening.

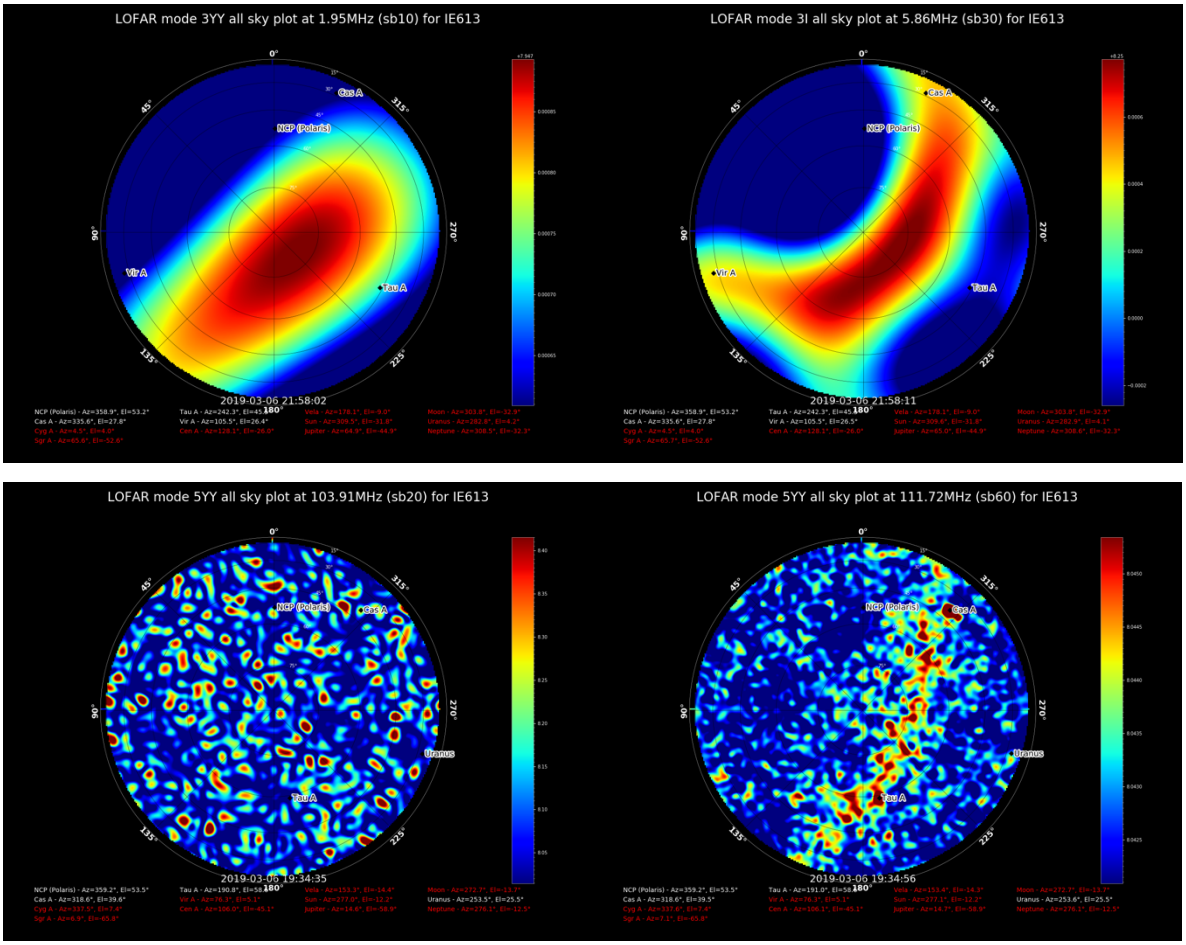

Figures 4.1, 4.2, 4.3, 4.4, All-sky plots of: mode 3 sub-band 10 (1.95 MHz, top-left), mode 3 sub-band 30 (5.86 MHz, top-right), mode 5 sub-band 20 (103.91 MHz, lower-left), mode 5 sub-band 60 (111.72 MHz, lower-right).

These images are typical of our mode 3 and mode 5 observations. From equation (1) we can note that angular resolution, ϑ, is dependent on wavelength λ and hence *inversely* proportional to frequency *f*. Increasing *f* decreases the telescope's resolution, allowing it to image smaller sources, but does not change the field of view. It is easy to conclude that imaging in the lowest sub-bands of mode 3 gives us too large of an angular resolution, resulting in only the most extended sources being resolved (namely, the Galactic Plane). Compare figures 4.1 and 4.2, separated by ten seconds but by 4 MHz; the shape of the source is already better-resolved in 4.2, extending to join the smaller but bright sources Cassiopeia A and Virgo A.

Meanwhile, imaging in mode 5 gives us quite a high angular resolution. Compare figure 4.2 to 4.4; the Galactic Plane is visible in both, but is resolved into a large collection of small sources, including Cas A quite intensely. This in itself is not a problem, but the low sub-band on mode 5 are clearly contaminated by some RFI due to being very near to the FM radio

band. Figure 4.3 in particular is unrecognisable as a sky source due to the strong RFI, but moving from its 103 MHz up to 111 MHz already lessens this effect. Repeated observations in the same band could also be used to eliminate the transient RFI contamination.

It is clear to see here two regions which we can rule out for all-sky imaging. Mode 3 images would be better taken in a higher sub-band, e.g. the 41 MHz band shown in the sample data in section 3.3. Mode 5 images likewise would be better taken in a range further from the FM edge of the band. We can conclude that while it is possible to make high-quality radio sky maps at Birr, careful consideration must be made of the chosen observation frequency and the sub-bands imaged within that.

4.2 Pulsar plots

The pulsar that we chose to image was PSR B0329+54, a very bright pulsar with S_{400} = 1500 mJy. We observed for two minutes at a time in each of receiver mode 3, 5 and 7. We processed lane 2 of each observation, so the plots below represent observations across the bands 30 – 40 MHz, 135 – 145 MHz, and 230 – 240 MHz. They are output generated by the `prepfold` command, which folds the data given to it, meaning it analyses the signal in search of a regular repeating pulse.

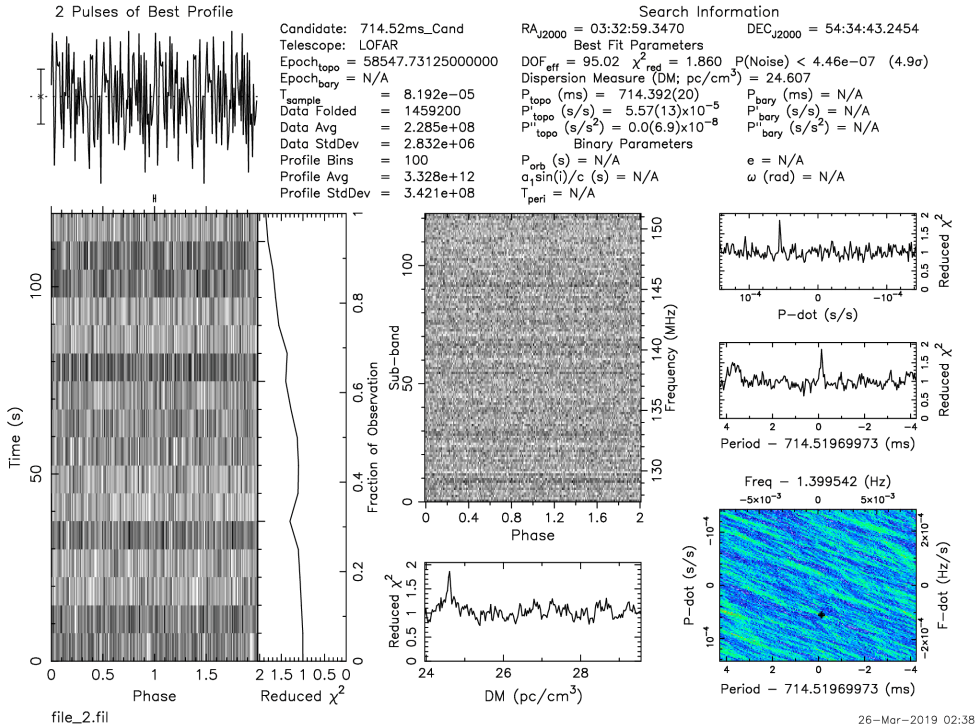

Figure 4.5: PRESTO plot of mode 3 observation of PSRB0349+54

This first plot displays our data from mode 3. We suspect that the pulsar was not bright enough at this low frequency (30-40 MHz) to obtain a clear observation of the pulse; no pulses can be discerned from the data in any of the plots above, even after supplying the known period and dispersion measure for the pulsar. We can see that the correct period was found (second graph on the right side), though not to a high degree of confidence at all. More care will need to be taken to attempt to get a clear, clean observation of a pulsar in this kind of frequency range.

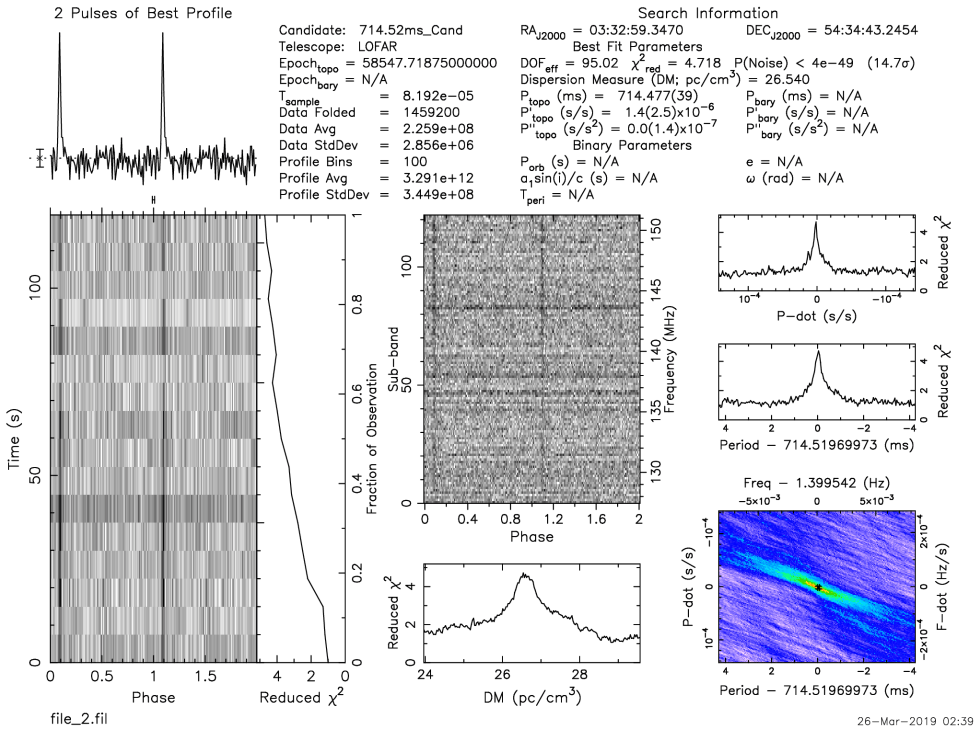

Figure 4.6: PRESTO plot of mode 5 observation of PSRB0349+54

This second plot displays our data taken in mode 5, and it is much more typical of a PRESTO output plot.

To explain the various data: the values along the top are observation, data, and search parameters used or found by PRESTO. The graph on the left side is a plot of time (120 seconds) against phase (rotation position of the pulsar, binned into smaller phase bins). While quite coarse due to our short observation, a clear black line is visible at two locations in the rotation. The data is shifted to attempt to show two pulses, and these align at the same location in each individual rotation, suggesting strongly that this is the pulse. Atop this is a simple intensity against phase graph which leads to the same conclusion. Alongside it is

a measure of the chi-squared value for this data, a statistical test standard for pulsars [e.g., Vaughan *et al.* 1994].

In the centre is a similar graph but for frequency against rotation phase, and here we see the same periodic signal appear again. The synchrotron radiation of pulsars is broad-spectrum, so we should expect it to cover the whole sub-band here, and it does so. Below is the results for PRESTO's search for the dispersion measure, with a peak at the value it is most likely to be according to the chi-squared test. Right of this again are similar graphs for the searches for period and period derivative (spin-down rate).

The last plot is a two-dimensional plot of period (and spin frequency) against period derivative (and frequency derivative). The colour on the plot represents likelihood here. The two values align in the centre of the plot, which is as expected for a pulsar since the values are linked, and confidence is low in other regions (although we can see it is more restricted for P-dot than it is for P). This graph gives us confidence that the data we have taken is valid and valuable data – see a similar good plot on Ransom [2015].

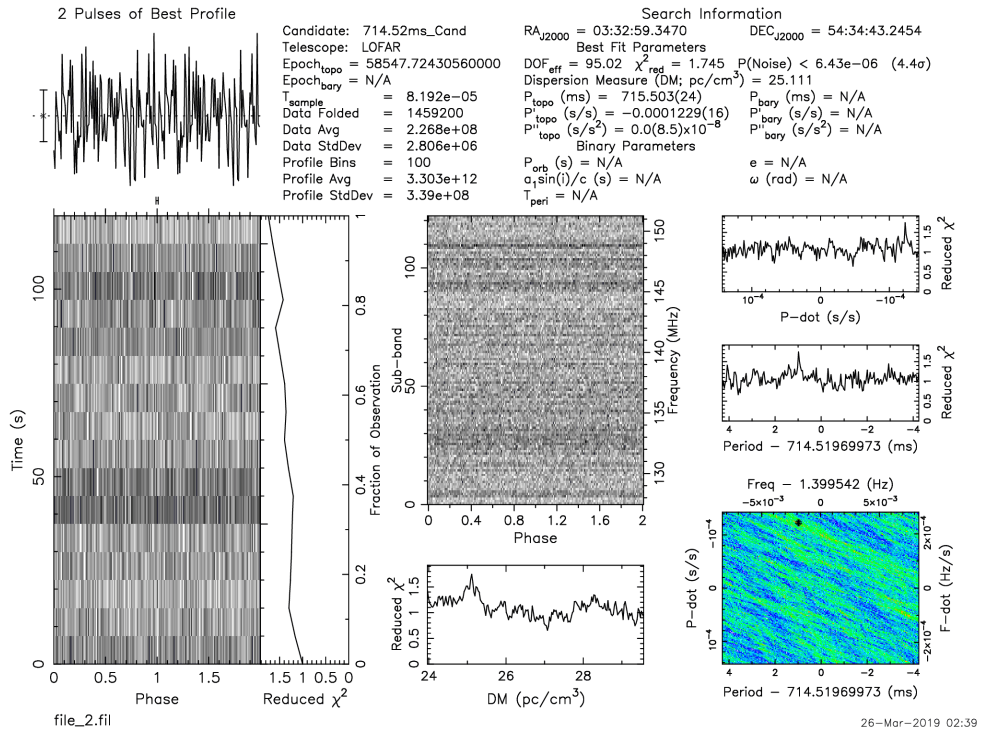

Figure 4.7: PRESTO plot of mode 7 observation of PSRB0349+54

This final plot is of our mode 7 observation. We expected to obtain a similar if not better plot than for mode 5 here, due to the higher flux density at higher frequency, but the search has drawn no conclusions, even with period and DM provided. We suspect that this

is therefore a processing error, either in the filterbanking or the analysis process, but a few attempts to reprocess the data failed similarly. While further efforts would be likely to uncover the error and obtain useful output plots, we note that this shows the current pipeline is far from perfect, and not yet fully-developed.

5 Conclusions and future work

Through our work with I-LOFAR and REALTA in this project, we were able to use the telescope and take data in two distinct observation types, as well as process and analyse it. We first note that this is a significant verification of the usability of the telescope and the feasibility of the methods in question. This was a relatively short ten-week project performed by an undergraduate student, and our success suggests that it is feasible for future students to utilise the telescope for introductory research. Having a research-grade telescope with modern technology accessible to students is a rare opportunity and our work suggests that there is much space here to make use of it.

Our all-sky imaging work re-demonstrated the ability of I-LOFAR to quickly and simply create all-sky radio maps. The difficulties in our analysis stemmed essentially just from the frequency ranges we selected, teaching us that these regions are not very suitable for analysis, and future work should avoid them. We suggest that an area of research could include mapping and verifying the frequency ranges available to I-LOFAR, so as to determine others that are unsuitable for mapping. The process of taking the data and processing it into maps is at this stage mature and our project shows that it can be done with relative ease.

An important aspect of the project was an investigation of the REALTA cluster's capabilities and methods, since it is such a new system and directly-accessible to researchers at NUIG. We were able to show that the system is already capable of basic detection of pulsars. While at current we cannot search data for unknown periods and dispersion measures of pulsars, we can verify them in data we take, and obtain period-derivative and other statistics. Further work in this area could focus on improving the analysis of data, which is likely to require just testing more data with PRESTO, until entirely unknown quantities can be detected and measured – it is in principle possible, just not yet with LOFAR data. We suggest that in the future I-LOFAR could be capable of monitoring known pulsars in this way, allowing it to search for 'glitches' and other variations, which is a

largely unexamined area of pulsar astronomy. Our work in this project has helped to further develop the pulsar-measuring capabilities of I-LOFAR, and we hope that it can be built upon.

6 Bibliography and references

6.1 Books used

Beskin et al., 1993. *Physics of the pulsar magnetosphere*. Cambridge, Canada ; New York: Cambridge University Press.

Burke, B.F. & Graham-Smith, F., 2010. *An introduction to radio astronomy*. 3rd edition, Cambridge ; New York: Cambridge University Press.

Lorimer, D.R. & Kramer, M., 2005. *Handbook of pulsar astronomy*. Cambridge, UK ; New York: Cambridge University Press.

Lyne, A.G. & Graham-Smith, F., 1998. *Pulsar astronomy*. 2nd edition, Cambridge ; New York: Cambridge University Press.

Thompson et al., 2017. *Interferometry and synthesis in radio astronomy*. 3rd edition, Cham: Springer International Publishing AG.

6.2 Reference list

Birkinshaw, M. and Lancaster, K. (2004) *Observational issues in radiometric and interferometric detection and analysis of the Sunyaev-Zel'dovich effects*, arXiv:astro-ph/0410336. Available: https://arxiv.org/abs/astro-ph/0410336 [accessed 4 April 2019]

I-LOFAR (2017) *Building I-LOFAR*. Available: http://lofar.ie/building-i-lofar/#2010 [accessed 2 April 2019]

I-LOFAR (2018a) *Big data processing for I-LOFAR: REALTA*. Available: http://lofar.ie/realta-launch/ [accessed 4 April 2019]

I-LOFAR (2018b) *First pulsar light by I-LOFAR*. Available: http://lofar.ie/first-pulsar-light/ [accessed 4 April 2019]

Isella, A. (2011) *Interferometry basics*, California Institute of Technology. Available: https://science.nrao.edu/opportunities/courses/casa-caltech-winter2012/Program [accessed 4 April 2019]

Lorimer, D. R. (2008) *SIGPROC – pulsar signal processing programs*. Available: http://sigproc.sourceforge.net/ [accessed 5 April 2019]

Manchester, R. N. *et al.* (2005) *The ATNF pulsar catalogue*, ApJ, 129 (5): 1993-2006. Available: http://www.atnf.csiro.au/research/pulsar/psrcat/ [accessed 3 April 2019]

McKay, D. (2018) *International LOFAR station hardware*, University of Tromsø KAIRA-UIT-PRS-507. Available: http://lofar.ie/i-lofar-school/ [accessed 15 January 2019]

McKenna, D. (2018) *allSkyImaging*. Available: https://github.com/David-McKenna/allSkyImaging [accessed 4 April 2019]

O'Brien, T. (2009) *Frontiers of modern astronomy*, University of Manchester Jodrell Bank Observatory. Available: http://www.jb.man.ac.uk/distance/frontiers/pulsars/section2.html [accessed 3 April 2019]

Ransom, S. (2015) *PRESTO*, National Radio Astronomy Observatory. Available: https://www.cv.nrao.edu/~sransom/presto/ [accessed 5 April 2019]

Vaughan B. A. *et al*. (1994) *Searches for millisecond pulsations in low-mass X-ray binaries*, ApJ, 435 (1): 362-371. Available: http://adsabs.harvard.edu/full/1994ApJ...435..362V [accessed 5 April 2019]

Virtanen, I. I. (2018) *Station data cookbook*. Technical Report LOFAR-ASTRON-MAN-064, ASTRON.

7 Acknowledgements

This project was undertaken with the support and under the supervision of Dr. Matt Redman, whom I thank sincerely for his sound advice and guidance across all aspects of the project technical and procedural, his constant encouragement and his generosity with his time.

I would like to thank also especially Joe McCauley for the time and effort spent assisting me in actually taking the observations which were essential to my project, providing me with his pulsar analysis scripts to use and helping me to understand the processing of high-frequency pulsar data, as well as generally shunting data around for me on the remote system.

I appreciate the help provided by Nevenoe Guegan throughout the project with some technical aspects as well as his availability to dispense advice; similarly I'd like to thank James Nallen and Derek Coburn for their advice and expertise given.

Additionally I am grateful for the cooperation of Peter Gallagher, Paul Callanan and his project student David Cleary, and Pearse Murphy in giving us advice on various aspects of the project and lending their time to do so.

Finally, I would like to acknowledge David McKenna, who made his excellent all-sky imaging scripts from his own project freely available with very comprehensive documentation and which I made extensive use of, and some modifications to, on that part of this project.

A. Appendix: Pulsar altitude and flux plotting

In preparation for investigation of pulsar emission as part of this project, it was necessary to select *which* pulsars to observe. This was done based on two criteria: a) whether they would be visible in the sky at the desired observing time; b) whether they would be bright enough at I-LOFAR's low observing frequencies to image. The object parameters and data used in order to make these choices were taken from the ATNF Pulsar Catalogue, PSRCAT [Manchester *et al*. 2005].

Two Python scripts were created for this purpose. One, using the `astropy` library, was used to plot the altitude of the targets for a given time and observing location. This was necessary because it is possible with I-LOFAR to observe during the day as the Sun does not usually interfere with radio observations as it does in optical. However, existing tools such as Staralt were mostly restricted to plotting only during night times. We decided it was easiest to create a similar tool to do so ourselves.

The second script was used to plot the flux density, *S*, of the pulsar beam with frequency. The reason for this was that PSRCAT only had data for flux at 400 MHz and I-LOFAR observes in the 10 – 270 MHz band. By determining a rule for the variation in flux density with frequency, we could propagate the S400 figure backwards to 150 MHz and 50 MHz. This was used to check if the pulsars were visible to us or not even if at high altitude.

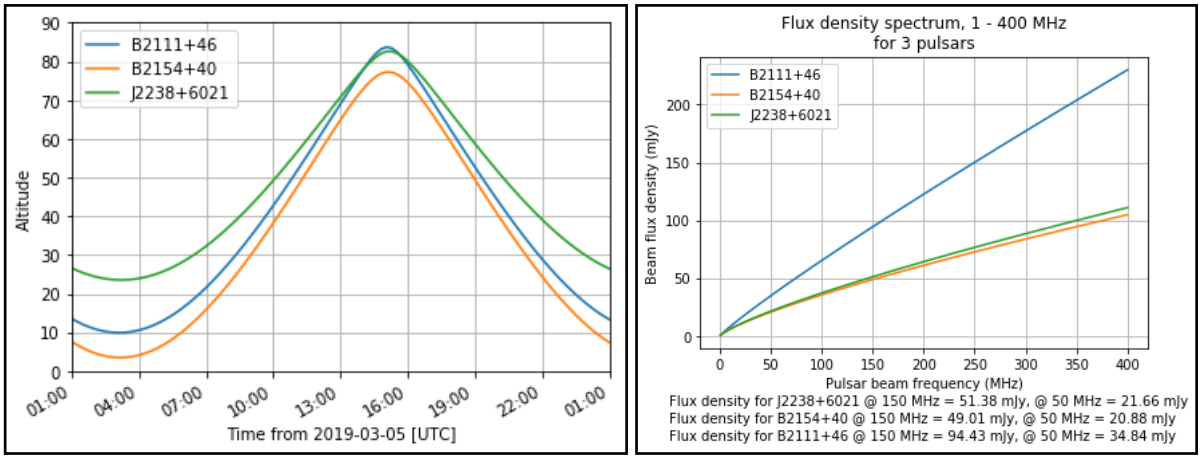

Figure A.1: plot of altitude (degrees above horizon) against time; A.2: of flux density against beam frequency. These plot initial candidate targets, not used because the observing time was later changed.

The flux script was based upon a power law equation of the form

$$S_f \propto f^{\alpha} \quad (A1)$$

from Beskin *et al.* (p. 25).

S_f is the flux density at emission frequency f; α is the spectral index of the pulsar, which determines how strongly its intensity changes with frequency. It is observed per O'Brien [2009] that intensity has a maximum value at around 400 MHz and decreases either side of this. Hence our graph plots only the decreasing slope behind that point, projected back to 0 MHz.

In the script the value α was calculated simply by reversing equation A1 and inputting the known S400 value like so

$$\alpha = \frac{\log(S_{400})}{\log(400)} \quad (A2)$$

where 400 is the value in MHz of f. α comes out as a negative value giving a negative slope as we expect and ranges typically between 2 and 4. This was then used to make the plot as in Figure A.2 above.

It is worth noting that the values obtained are approximate as the true value of S_f is only proportional, not equal, to f^{α}. For our purposes this was sufficient to give us an idea of whether or not the pulsar would be possible to image; the precise flux values were unimportant. We selected pulsars which were at least still in the tens of mJy of flux density at S_{50} for imaging.