New Instruments in Low Frequency Radio Interferometry and their Calibration: LOFAR and MWA

Nicholas LeBel-Buchanan

Master of Science

Physics

McGill University

Montreal,Quebec

2013-12-13

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of Master of Science

DEDICATION

This thesis is dedicated to my wife, without whom I never would have made it.

ACKNOWLEDGEMENTS

I thank my supervisor Matt Dobbs for his dedication and support throughout this project. I also thank Keith Vanderlinde and Kevin Bandura for their frequent help, and Ivan Padilla for many fruitful discussions.

ABSTRACT

Low frequency radio astronomy is experiencing a renaissance due to advances in computing power and efficiency. New instruments are being designed and built to take advantage of this. An introduction is given to radio astronomy. The physical mechanisms of radio emission are explained. The measurement of radio waves is discussed, particularly in regard to the properties and advantages of interferometric arrays. Two modern instruments currently being deployed and tested are described, with particular attention paid to their calibration techniques. These instruments are the Murchison Widefield Array (MWA) in Western Australia and the Low Frequency Array (LOFAR) in Europe, with most components in The Netherlands. These instruments are compared with one another.

ABRÉGÉ

La radioastronomie de fréquence basse connaît une renaissance grâce aux progrès de la puissance de calcul et d'efficacité. De nouveaux instruments sont conçus et construits pour profiter de cette. Une introduction est donnée à la radioastronomie. Les mécanismes physiques de l'émission de radio sont expliqué. La mesure des ondes radio est discutée, en particulier en ce qui concerne les propriétés et avantages de réseaux interférométriques. Deux instruments modernes actuellement déployés et testés sont décrits, avec une attention particulière à leurs techniques d'étalonnage. Ces instruments sont le "Murchison Widefield Array" (MWA) en l' Ouest de l'Australie et le "Low Frequency Array" (LOFAR) en Europe, la plupart des composants dans les Pays-Bas. Ces instruments sont comparé l'un à l'autre.

TABLE OF CONTENTS

LIST OF FIGURES

Figure page

CHAPTER 1
Introduction

The discipline of radio astronomy is only about eighty years old. It began when workers at Bell Labs first noticed an unexpected background hiss in their radio equipment that was ultimately determined to be of extraplanetary origin. The first radio map of a large region of the sky was first conducted in 1941 on an amateur-built 9 m dish. Since then, radio telescopes have grown larger, more complex, and more sensitive. They have led to the first observations of astrophysical synchrotron radiation, the discovery of pulsars and quasars, and the measurement of rotation curves of distant galaxies. The use of interferometers such as the Very Large Array (VLA) and Very Long Baseline Interferometry (VLBI) have enabled extremely high resolution imaging. The largest long baseline array, a combination of the European and American VLBI instruments called the global VLBI, is capable of sub microarcsecond resolution at high frequency [1].

In recent years, computing power has reached the point that so-called "software telescopes" are possible. These are generally composed of a large number (hundreds or thousands) of fixed receiving elements. The signals from these are combined by a computer and using phase shifts, delays, and other techniques the array can be 'pointed' after digitization. This has allowed tremendous simplification of the mechanical hardware. As computing power continues to follow Moore's

Law these telescopes will become more common and more powerful. New algorithms for calibration and data reduction must be developed to make full use of the flood of data now available.

This thesis will give an introduction to the science and technology of radio astronomy and interferometry. Then with a view to understanding these new challenges and capabilities, two modern instruments will be analyzed: the LOw Frequency Array (LOFAR) base in The Netherlands and the Murchison Widefield Array (MWA) in the Australian Outback. Particular attention will be given to the calibration of these instruments. The size and sensitivity of these instruments present new challenges in this area [2, 3].

CHAPTER 2
Introduction to Radio Astronomy

Radio astronomy refers to the study of radio emission from any non-man-made source outside the solar system. The basic physical principles of radio emission and measurement are presented in section 2.1. Radio telescopes, particularly interferometers, are discussed in section 2.2.

2.1 Physical Processes

This section presents the basic principles of radio waves and their interactions. The physical mechanisms of radio emission are explained and examples of emitting sources given. The measurement of radio waves is also discussed.

2.1.1 Electromagnetic Waves

The electromagnetic spectrum is divided into many bands. The distinctions between these bands are arbitrary; in terms of the waves themselves, they bear no physical meaning. They exist as they do largely owing to the history of their discovery.[4] The radio portion of the spectrum spans wavelength 10^{-1} m $\leq \lambda < \infty$ or, equivalently, frequency 3×10^9 Hz $\geq \nu > 0$ [5].

Electromagnetic (EM) radiation can have two possible polarizations. There are two common bases used to describe the polarization state: linear or circular polarization. In the linear view, two perpendicular axes are arbitrarily selected and the EM wave is described by the amplitude of electric (or magnetic) field oscillations

along each axis. There is also some relative phase between the two polarizations. Circular polarization is slightly harder to visualize. It also has two basis states, called left and right hand circular polarizations, which also have some relative phase between them. A wave with a pure circular polarization has an electric (or magnetic) field vector which retains a constant magnitude but rotates at the frequency of the wave. Right handed waves rotate in the sense of a corkscrew, left handed waves the other way. Both these representations, if the relative phase is known, completely describe the EM field of a wave. As is shown below, it is possible to convert from one basis to the other. Let E be the magnitude of the electric field and the subscripts x, y, l, and r refer to the x-axis, y-axis, left, and right polarizations respectively. The following relations hold:[4]

$$E_x = \frac{1}{\sqrt{2}}(E_l + E_r) \quad E_y = \frac{1}{\sqrt{2}}(E_l - E_r)$$
$$E_l = \frac{1}{2}(E_x + E_y) \quad E_r = \frac{1}{2}(E_x - E_y)$$

2.1.2 Radiative Mechanisms

Several naturally occurring emission mechanisms exist which produce radio waves. Common sources include thermal, synchrotron, transition, and magnetic dipole radiation.

All macroscopic objects emit thermally to some extent. It should be noted that the classic blackbody spectrum is an approximation which breaks down for wavelengths larger than the emitting object. Stars, for example, emit radio waves by this mechanism. Dust, which is much smaller than any radio wave, does not.

Synchrotron radiation is emitted by high energy electrons travelling in a magnetic field. Because of the force exerted on a moving charge by a magnetic field, electrons spiral around magnetic field lines. The period of the spiral is dependent on both the momentum of the electron and the strength of the magnetic field. Since an accelerated charge radiates, these electrons emit at a frequency in accord with their period of motion. This radiation is linearly polarized in the plane perpendicular to the magnetic field [6, 5].

Atomic transitions can radiate in the radio frequency. In astronomy, the ionization states of atoms are indicated by their chemical symbol followed by a Roman numeral. A chemical element "X" will be written XI for neutral X, XII for singly ionized, XIII for doubly ionized and so on. This applies to individual atoms, not to molecules. For example, neutral atomic hydrogen (HI) has an emission line at 1.4 GHz. This occurs when the spins of the electron and proton transition from aligned to anti-aligned. The larger part of the gas in any galaxy is HI so this has been used to map their matter and velocity distributions.

If a magnetic dipole varies with time, it will radiate. This is precisely analogous to the electric dipoles common in communication equipment. Pulsars have extremely strong magnetic fields which oscillate on time scales of ~ 1 second or less, and this is what gives rise to the radio pulses we measure. The interaction of the magnetic field with the material of the neutron star is complex and the exact emission mechanism of the waves we see is disputed. The ultimate driver however is known to be the oscillating magnetic field [5].

2.1.3 Measurement

Antennas are used in radio measurements. Antennas couple directly to either the magnetic or electric field of a propagating wave, rather than to the power, as in a bolometer or CCD. It is most common to use a dipole antenna or some derivative thereof to couple to the electric field. Circularly polarized antennas are rare in astronomy. Though there is some variance from one design to another, dipoles are approximately half the wavelength they are designed to measure. Magnetic loop antennas, consisting of coils which couple to magnetic fields in much the same manner as a transformer, are usually only used when a dipole would be inconveniently large. Magnetic antennas lose sensitivity less severely when they are far from their optimum size. The ionosphere is reflective at wavelengths over 10 meters, so any such waves detected are certainly of terrestrial origin [5]. This is the only region of the radio spectrum where magnetic loop antennas are commonly used, so they will not be discussed further. Dipoles are often used in collocated perpendicular pairs to provide sensitivity to both polarizations. An antenna is also referred to as a feed.

Radio telescopes often use a large reflector to concentrate incoming radiation and give directional sensitivity. Reflectors are usually parabolic to preserve coherence and direct all energy onto a pair of crossed dipoles. This is illustrated in figure 2–1 on the next page. This is analogous to a single pixel camera. Some instruments such as the Parkes Multibeam Receiver have used multiple feeds to provide several simultaneous observations.

The sensitivity of an antenna is a function of direction and wavelength. At any particular wavelength, this function is called the beam pattern of the antenna.

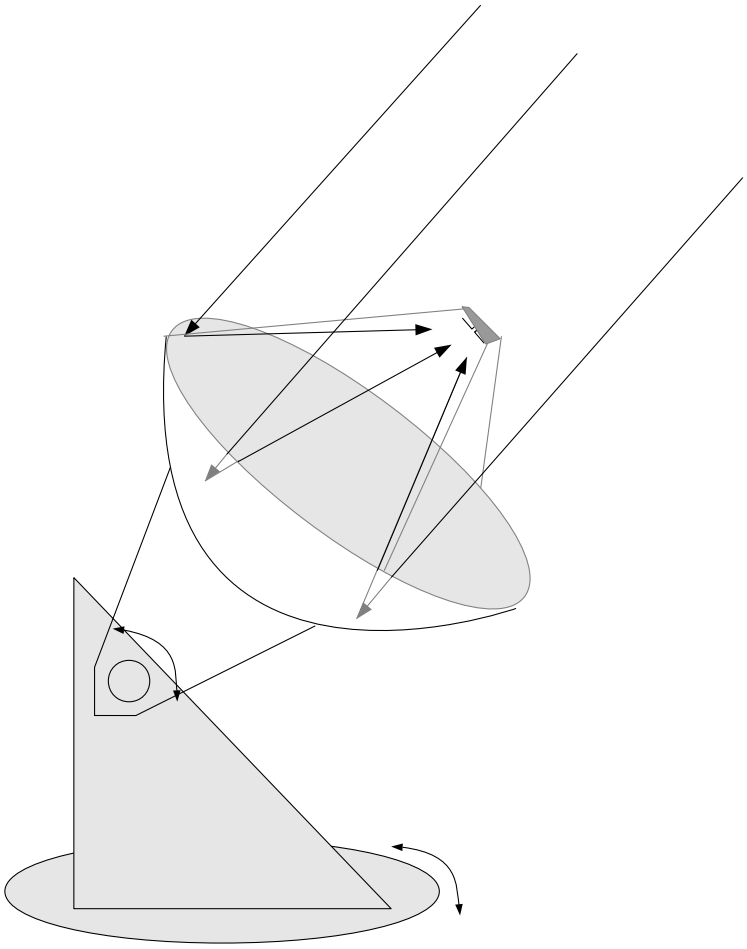

Figure 2–1: A common radio telescope configuration is a large parabolic reflector which focuses the incoming signal onto a single crossed dipole. As this makes a single measurement of the sky in a single direction, multiple measurements are necessary to build up a map. Studies of point sources do not need this. Radio telescopes are often steerable in one or both axes which can allow scanning or tracking of a source.

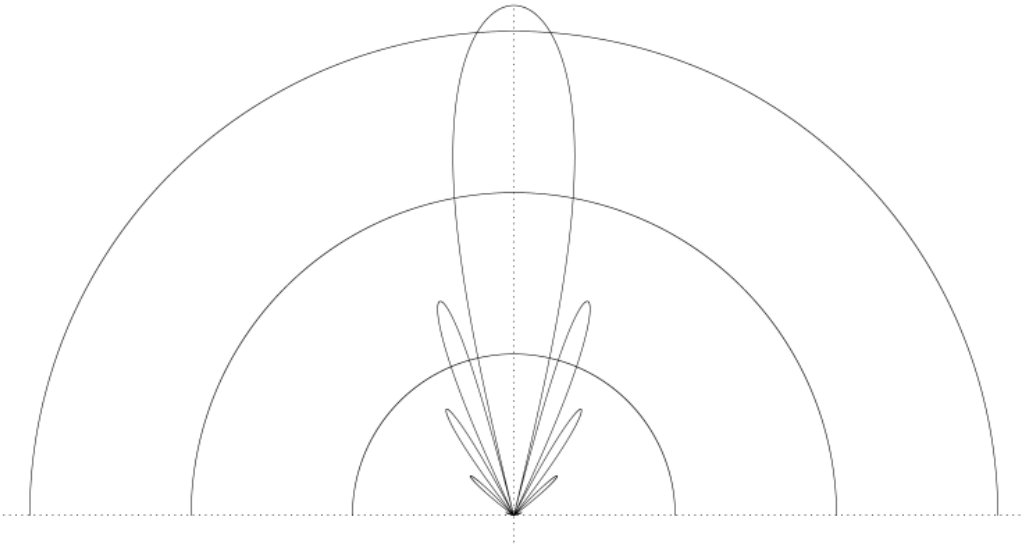

Figure 2–2: Idealized forward beam pattern of a parabolic reflector. This assumes even illumination of the dish, which holds well in a plane perpendicular to the axis of a dipole. Distortions from dish surface imperfections, the shadow of the feed, and struts have been ignored. The dish diameter is taken to be 5λ. Circles at -10 dB, 0 dB, and +10 dB gain are shown. This gain is relative to an isotropic feed. The gain of this reflector is 11.6 dB.

Elements of a telescope such as the reflector or the struts which support the receiver also interact with the incoming waves and modify the beam. Indeed, that is the purpose of the reflector. Beams are often sufficiently azimuthally symmetric that they may be adequately represented in two dimensions. Figure 2–2 shows the idealized forward beam pattern for a simple dipole mounted at the focus of a parabolic dish [7]. This simulation ignores the effect of the struts supporting the feed and imperfections in the dish.

Radio telescopes use very large collecting areas for two reasons. One is that the spectral flux density from radio sources is extremely low. The unit of measure

used for point sources is the Jansky: $1\ \text{Jy} = 10^{-26}\frac{\text{W}}{\text{m}^2\text{Hz}}$. This is a typical value. Large reflectors gather more of this power into the receiver. The other great boon of large collecting areas is to increase resolution. Rayleigh's criterion states that the maximum resolving power for any circular aperture telescope is $\theta_{min} = 1.22\frac{\lambda}{D}$ where:[5]

θ_{min}	is the minimum angular separation between resolvable point sources in radians
λ	is the central observing wavelength
D	is the diameter of the aperture (the reflector dish in this case)

For example, a radio telescope operating at 1.4 GHz with a 10 m dish, which would be extremely large for an optical telescope, would have a resolution of only 1.5º. Increased resolution is valuable not only for locating a source on the sky, but also for detecting faint sources against a diffuse background.

2.1.4 Temperatures and Effective Temperatures

The ubiquity of blackbody radiation makes it natural to consider the temperature of an observed source. If a compact source is resolved (such as Sol) then the physical temperature relates directly to the power received as shown in equation 2.1. This relies on the Rayleigh-Jeans approximation.[8]

$$P = kT\Delta\nu \tag{2.1}$$

Unresolved point sources are said to have an effective temperature T_{eff}. This effective temperature is defined as the temperature a resolved source would need to

have to induce the same response in the receiver. The total radio power detected by a receiver with bandwidth $\Delta\nu$ centred at ν and collecting area A from a point source with spectral flux density $S(\nu)$ and resolved source of temperature T is given in equations 2.2 and 2.1 [5].

$$P_{\text{pt. src.}} = A \int S(\nu)\,\mathrm{d}\nu \tag{2.2}$$

Setting the above powers equal and solving for temperature leads to a simple definition of effective temperature.

$$T_{\text{eff}} = \frac{A}{k\Delta\nu} \int S(\nu)\,\mathrm{d}\nu \tag{2.3}$$

Effective temperature of a point source is instrument dependent and should not be construed as having physical significance.

Diffuse sources are often not blackbodies. Synchrotron radiation is an example of this [5]. Thus even though they are fully resolved, the power received does not relate directly to their physical temperature. Nonetheless, they may be quoted as having an effective temperature using equation 2.3. This temperature will be a function of frequency.

Noise, especially amplifier noise, is also frequently quoted as a temperature using a form of equation 2.1. Temperature is made equivalent to noise power by considering a resistor connected to the input of a hypothetical perfect (noiseless) amplifier. The temperature of the resistor that would generate the same electrical

noise power as the real amplifier does is said to be the temperature of the amplifier:

$$T_{\text{amp}} = \frac{P_{\text{noise}}}{k\Delta\nu}$$

The other significant noise contributor is ground pickup, also called spillover. Some of the sidelobes of an antenna beam will be pointed at the ground or other warm objects, which emit blackbody radiation. A very small proportion of the sensitivity of a well designed instrument will be directed at the ground, but in general the ground is much hotter than any source on the sky so it can still be significant. The spillover temperature is the product of the physical temperature of the ground and the fraction of the sensitivity directed at it.

While not all of these effective temperatures are physical, they are directly comparable. They are all measurements of received power, albeit in peculiar units. They can thus be used to compare sources of different kinds, or estimate integration times that may be necessary to detect some particular source, or, since some of these are instrument dependent, choose which instrument may be suitable for a particular observation.

2.2 Interferometry

Interferometry takes advantage of the same phenomenon observed in Young's double slit experiment. Any coherent combination of signals from multiple antennas is interferometry, whether the signals are added or correlated, in analog or digital, in real time or in post-processing. All of these variations are in use in modern instruments. Several significant gains are made from this coherent approach to a multiple antenna system. The resolution is greatly increased: the aperture size is

no longer the size of the dishes, but the distance between them. This is of particular importance in studying point sources because of the greatly reduced sky noise in a small beam. Interferometers also have multiple beams on the sky simultaneously, increasing mapping speed in proportion to the number of antennas.

2.2.1 Early Two Element Interferometers

The first interferometers consisted of two receiving elements, consisting of antennas with reflectors or dipole arrays, whose signals are brought together and summed before detection. This situation is illustrated in figure 2–3 on the next page.

If a source is located at some angle θ from zenith, there will be a delay in the arrival of a wavefront at one antenna as compared with the other. This delay we call τ. We will consider the simple case of a narrow band receiver centred at frequency ν_0 so that the received wave may be taken to be sinusoidal. Let D be the distance between the two antennas and d the length of retardation of a wavefront at one antenna relative to the other.

$$d = D\sin(\theta)$$

$$\tau = \frac{d}{c} = \frac{D\sin(\theta)}{c}$$

$$V_1(t) = V_0\sin(2\pi\nu_0(t-\tau))$$

$$V_2(t) = V_0\sin(2\pi\nu_0 t)$$

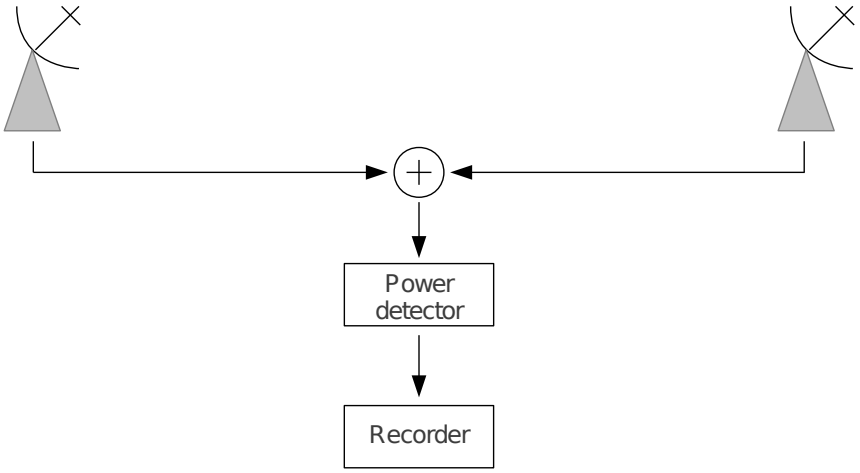

Figure 2–3: Schematic of a 2-element interferometer. The voltage signals from two separate antennas or telescopes are added by an analog circuit before any measurement is made. Early detectors were sensitive to the input electrical power, not the voltage. Phase information was lost.

The detector in radio telescopes is usually sensitive to power, so in this case it will be detecting the power in the sum of the two voltages: $V_{\text{out}} = V_1 + V_2$.

$$P_{\text{out}} = V_{\text{out}}^2 = (V_1 + V_2)^2$$

$$P_{\text{out}} = V_0^2 \left[\sin\left(2\pi\nu_0\left(t-\tau\right)\right) + \sin\left(2\pi\nu_0 t\right)\right]^2$$

The detector integrates the power over a time significantly longer than the radio frequency being detected, and thus functions as a lowpass filter. Eliminating terms in odd powers of $\sin(2\pi\nu_0 t)$, which integrate to zero, it can be shown that the above reduces to

$$P_{\text{out}} = V_0^2 \left[1 + \cos\left(2\pi\nu_0\tau\right)\right]. \tag{2.4}$$

This idealized, noiseless power may be taken as the portion of the response due only to the source, or equivalently as the beam of the interferometer configuration. It is clear that τ is a function of the Earth's position in space relative to the observed source, and so will slowly change with time as the Earth rotates. This will result in a signal fluctuating between 0 and $2V_0^2$. These are called interference fringes. An idealized example is shown in figure 2–4 on the following page.

If noise is taken into account, the above equation is modified to become equation 2.5.

$$P_{\text{out}} = N^2 + 2N\left(V_1 + V_2\right) + V_0^2\left[1 + \cos\left(2\pi\nu_0\tau\right)\right] \tag{2.5}$$

Where N is the voltage contributed by noise terms, such as the amplifier, ground pickup, RFI, backgrounds on the sky, etc. The noise is statistically uncorrelated with the signal, so the second term will tend to 0 with integration, but the autocorrelation

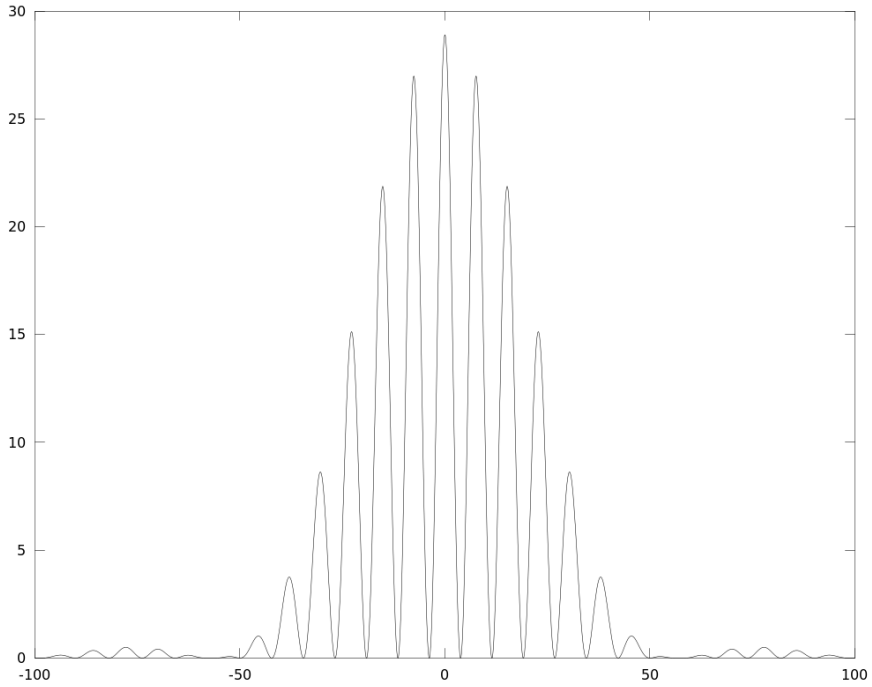

Figure 2–4: Interference fringes for a two element interferometer sensitive at $\nu =$ 300 MHz ($\lambda = 1$ m) with 5 m parabolic reflectors spaced at 30 m on an East-West line. These dishes would have the beam pattern shown in figure 2–2 on page 8. It has been assumed that the dishes do not track the source, but rather allow it to drift overhead, and that the source is at 0° declination. The horizontal axis is in minutes from zenith. The vertical axis is power in arbitrary units.

(N^2) of the noise will not. It is positive definite. In general the noise term can much larger than the signal term, and varies in magnitude. This proved to be a problem for chart recorders, which had small dynamic range and would often go off scale, for example with fluctuations in amplifier gain [8].

This problem was largely addressed when phase switching was introduced by Martin Ryle in 1952 [8]. In this system one of the signals periodically undergoes a phase shift of 180º, so that the output of the receiver is alternately $(V_1+V_2)^2$ and $(V_1-V_2)^2$. A second detection stage, which is synchronized with the phase shift, takes the difference between these two signals. Figure 2–5 illustrates this system. This serves to isolate the desired signal from the noise. Let N_1 and N_2 be the noise contributed by the two antennas.

$$P_{\text{out}} = (N_1+V_1+N_2+V_2)^2 - (N_1+V_1-N_2-V_2)^2 \tag{2.6}$$

It can be shown that this reduces to

$$P_{\text{out}} = 4V_1V_2 + 4\left[N_1N_2 + N_1V_2 + V_1N_2\right] \tag{2.7}$$

or

$$P_{\text{out}} = 2V_0^2\cos\left(2\pi\nu_0\tau\right) + 4\left[N_1N_2 + N_1V_2 + V_1N_2\right]. \tag{2.8}$$

The second term is composed entirely of products that are uncorrelated, and so will be reduced by longer integration times. Note that there is now no noise autocorrelation term. It was this positive definite term which caused the most significant

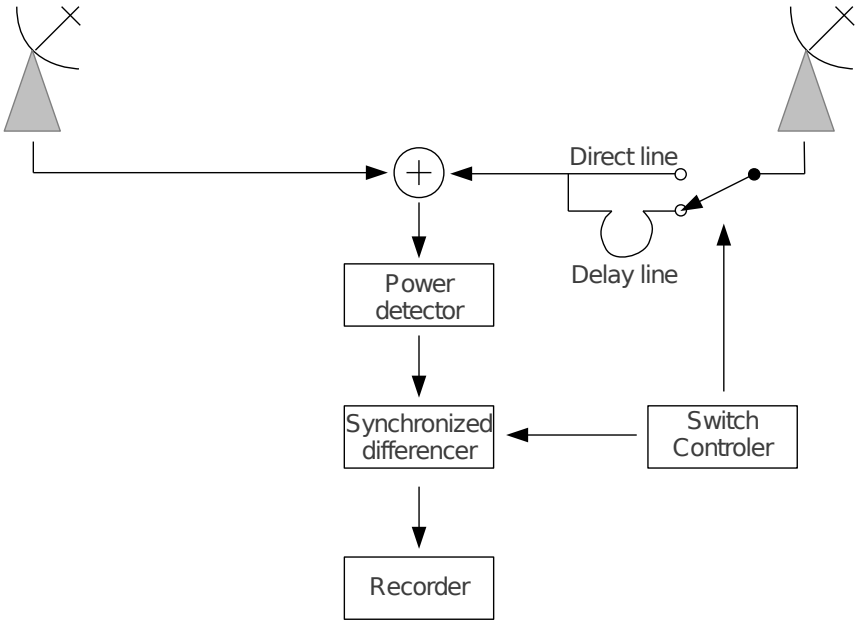

Figure 2–5: Phase switching correlator. By introducing a delay line of half a wavelength in length into one of the feed lines, the phase of this antenna is effectively inverted. This is equivalent to changing the addition of the voltages to a subtraction. A differencer synchronized with the delay line switch takes the difference between the power of the added and subtracted voltages. As is shown in equation 2.7 this produces a cross-correlation measurement.

problem in the early correlators. Without this generally large additive term, the measured correlated power will fluctuate about zero and the signal can be kept within the instrument's dynamic range easily.

2.2.2 U-V Space

In radio astronomy parlance, a cross-correlation between two antennas is called a visibility. An array may have many antenna pairs with the same spacing. In this case, each such pair is a separate measurement of the same visibility; the visibility is defined by the relative separation of the antennas, not their absolute location. Each visibility is the amplitude and phase of a Fourier mode on the sky. This mode is defined by the separation and orientation of the antenna pair, and its angular frequency is the fringe frequency of that antenna pair at a given wavelength .

Visibilities may be represented graphically in an inverse space called the u-v plane, where the u and v axes represent the separation between antennas in each direction. The most natural unit to measure this in is wavelengths. This plane is mapped onto the sky by a Fourier transform. The orientation of the axes with respect to the physical telescope is arbitrary, but it can be convenient to associate the u (v) axis with the east-west (north-south) separation. These separations are subject to projection effects along the line of sight; any given telescope will have different coverage in u-v space for different source directions. Figure 2–6 shows a simple array and the points in u-v space which it samples.

Some arrays have one or more movable antennas. Different configurations can have very different u-v coverages, which can be adjusted according to the science of

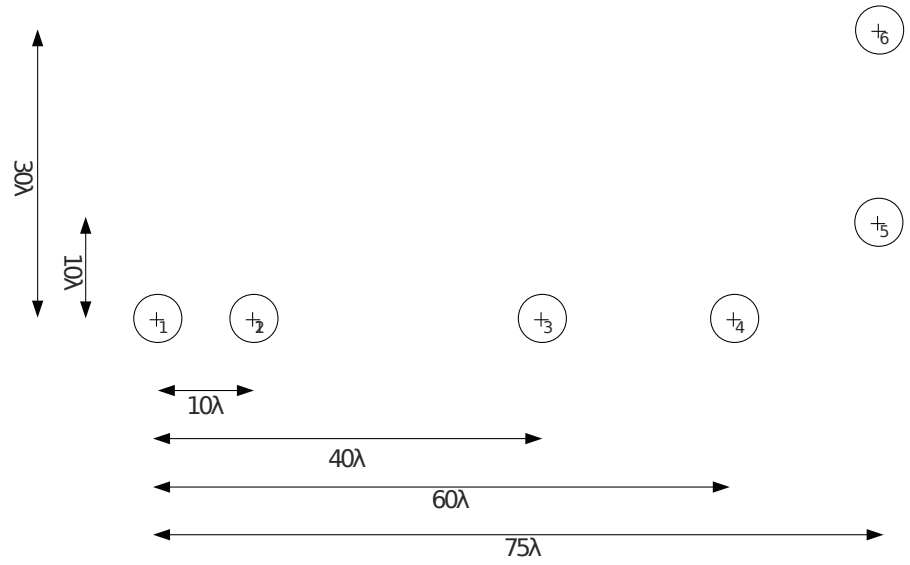

(a) Array map

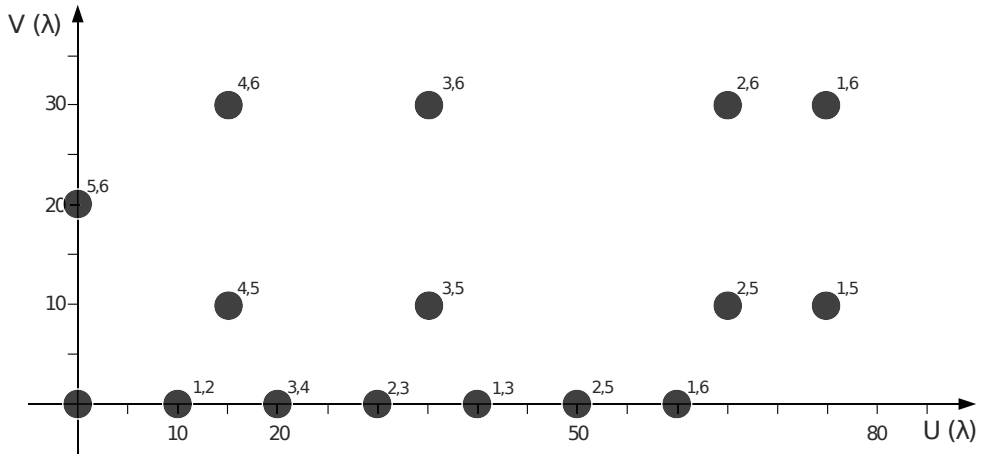

(b) u-v coverage

Figure 2–6: (a) Schematic map of an example array. The individual radio dishes are identified numerically. (b) U-V coverage of the array in (a). The antenna pair responsible for each point is indicated by the numbers beside it. The points have been given size to indicate that owing to the physical size of the dishes, any given cross-correlation is sensitive to some small range of scales.

a particular investigation. Data can be combined from observations taken with different configurations. This is referred to as aperture synthesis. It relies on the source under observation to have not changed significantly between the observations.

As noted above, u-v coverage is subject to projection effects along the line of sight. If a source is tracked, this projection changes as the Earth rotates. This can be used to fill in the coverage. This is called Earth rotation aperture synthesis.

In general there are many points in u-v space which are not sampled. These correspond to modes on the sky about which the instrument has no information. If these are naively assumed to be zero and a map is created by Fourier transforming the u-v plane, significant artifacts will be created. This is called a 'dirty image.' It accurately reflects all information which has been measured by the instrument. If there are reasonable assumptions that can be made or other information is available, (for example that the observed source is pointlike) then this information can be used to 'clean' the image.

2.2.3 Element Distribution

It is difficult to find, or even define, the optimum arrangement for a many-antenna interferometer. A variety of arrangements have been used, including X, T, and Y shaped arrangements, as well as linear arrays and quasi-random layouts. .

There are two types of array which deserve special attention. Large arrays with regular spacing can offer a computational advantage, such as using the Fast Fourier Transform (FFT) algorithm to form beams. Close-packed arrays which have antenna elements spaced no more than half a wavelength apart are a subset of this configuration. This Nyquist samples the incoming wavefront, and thus retains all

available information in the area of the telescope. The other case is a full-coverage array which produces a Nyquist sampled u-v space (up to some maximum baseline). This produces an image with no ambiguity up to the scales sampled by the interferometer, and no cleaning is necessary. Calibration, however, is still necessary.

2.2.4 Regularly Spaced Arrays and FFT Beamforming

Beamforming refers to the practice of combining signals from all the antennas into individual beams on the sky before the data is permanently recorded. The combined signal will have a response to the sky which is a function of the size and number of elements of the array. This is a formed beam, as distinct from the beam of an individual antenna. The manner in which the beam is formed is intended to optimize some desired feature of the response. This may be for example sensitivity in a particular direction, or low sidelobes. It is possible to form several beams simultaneously, up to the number of antennas. More beams can be created, but will be redundant. If there are N antennas, all the information is available in N orthogonal beams, where orthogonal should be taken to mean linearly independent, not that they are pointed in perpendicular directions.

Regularly spaced arrays result in a regularly sampled u-v plane. The spacing between every pair of antennas is some multiple of the smallest spacing. This is advantageous when converting the visibility data to real space with an inverse FFT. However, in this case it is not necessary to use visibilities at all: beamforming can be done very computationally efficiently and in principle the result contains all the information that would have been available in the visibilities.

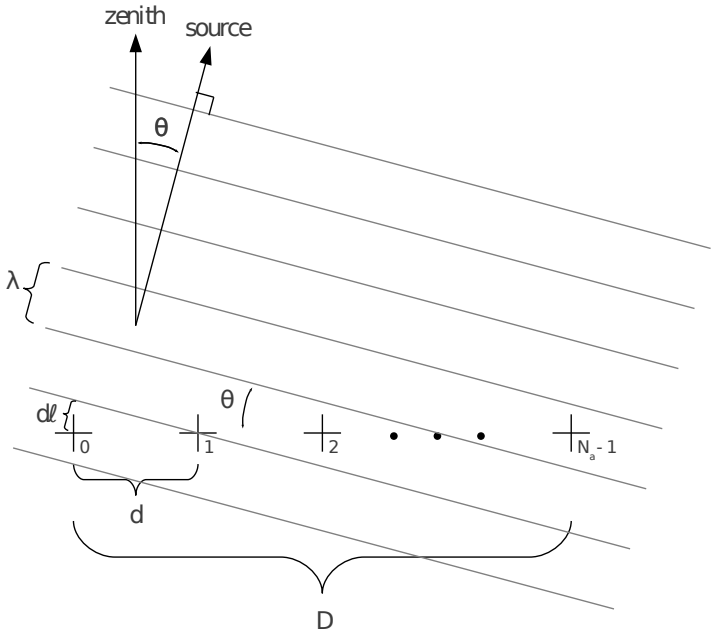

Figure 2–7: Regularly spaced antenna array with N elements. An incoming plane wave of wavelength λ which is incident at an angle θ will be retarded by a distance $d\ell$ from one antenna to the next. This is equivalent to a phase delay $\phi = 2\pi\frac{d\ell}{\lambda}$ or time delay of $t = \frac{d\ell}{c}$, where c is the speed of light.

Consider a linear array of N_a antennas set at a spacing d with bandwidth $\Delta\nu$ centred at ν_0. The arrangement is shown in figure 2–7. Each antenna will have some response $V_n(t)$. Let n range from 0 to $N_a - 1$. We will neglect noise terms in this analysis. The simplest means to combine the signals is to add them, which creates a beam with its maximum at zenith. At some incident angle θ a wavefront will be shifted between adjacent antennas by a distance $d\ell$. This is equivalent to a

phase shift of $\phi = 2\pi\frac{d\ell}{\lambda}$ or $\phi = 2\pi\frac{d\sin(\theta)}{\lambda}$. To create a beam directed at some angle θ from zenith we employ a phase shift to the signal from each antenna as follows:

$$B_\theta(t) = \sum_0^{N_a-1} V_n(t)\,\mathrm{e}^{-in\phi} \tag{2.9}$$

where $\phi = 2\pi\frac{\nu d\sin(\theta)}{\mathrm{c}}$. The Discrete Fourier Transform (DFT) of a series of values $\{x_0, x_1, \ldots, x_{N-1}\}$ is

$$X_k = \sum_{n=0}^{N-1} x_n \mathrm{e}^{-2\pi\imath kn/N}$$

Thus if we set $\frac{k}{N} = \frac{\nu d\sin(\theta)}{\mathrm{c}}$ one can see that $B_\theta(t)$ is a Fourier component of the sequence $\{V_1(t), V_2(t), \ldots, V_{N_a-1}(t)\}$. This allows the technique of FFT beamforming. At each timestep one Fourier transforms the sequence of voltages $\{V_1(t), V_2(t), \ldots, V_{N_a-1}(t)\}$. Each bin in the FFT output is the power in one of a fan of beams which are located at $\theta = \sin^{-1}\left(\frac{k\mathrm{c}}{N\nu d}\right)$ for k an integer from 0 to $N_a - 1$. This fan of beams is shown in figure 2–8 on the following page.

This form of data reduction maximally reduces the data rate without loss of information for an ideal telescope. For real telescopes there can be different beam patterns at each antenna and uncorrelated gain modulations of the pre-digitization amplifiers. The full N^2 correlation products contain not only information about the sky, but by comparisons with one another information about the instrument. This information can be used for calibration, so N^2 correlation is sometimes preferred even for regularly spaced arrays.

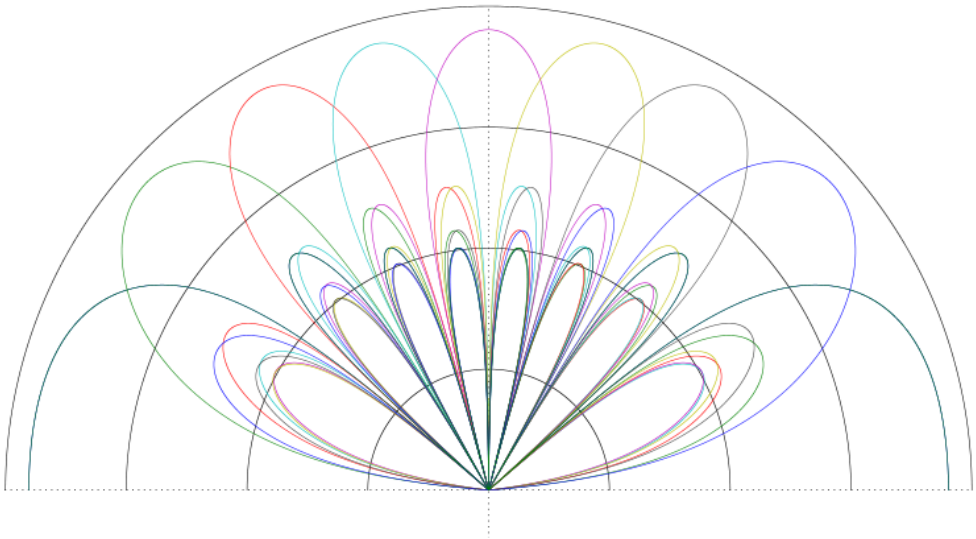

Figure 2–8: Beams formed through FFT beamforming. This is for an array with 8 antennas spaced $\frac{\lambda}{2}$ apart. The colours indicate each of the 8 beams. Note that the beam on each horizon is actually the same beam. The two horizons cannot by distinguished in this case. The arcs represent gradations of -10 dB, 0 dB, +10 dB, and +20 dB.

2.2.5 Full Coverage Arrays

Full coverage arrays are a special case of the quasi-random arrangement. They are designed to produced a fully sampled u-v plane. A regularly spaced array also produces an evenly sampled u-v plane out to the maximum baseline, but there are many duplicate small baselines. There will inevitably be either gaps in the u-v coverage or redundant baselines. The core of the Murchison Widefield Array is an example of a full coverage array. It will be discussed extensively in chapter 4.

CHAPTER 3
Low Frequency Array

3.1 Introduction to the Instrument

In this section an overview of the LOw Frequency ARray (LOFAR) is presented. LOFAR was conceived as a low frequency survey instrument. Early modifications to the design permitted Epoch of Reionization (EoR) science as well. Most information in this section is derived from "The LOFAR Telescope: System Architecture and Signal Processing" by Marco de Vos et al [9].

3.1.1 Design

LOFAR is an aperture synthesis instrument sensitive from 15 MHz to 240 MHz with substations all across Europe. The dense population makes this a very bad environment for Radio Frequency Interference (RFI). For this reason, the FM band (from 87 MHz to 108 MHz) is excluded. Two different antenna designs are used, one for high frequency and the other for low. Both are minimally directional active dipoles. The substations themselves are phased arrays. Processing at the substation level provides significant flexibility in terms of bandwidth and number of beams on the sky. There are 48 substations, 40 of which are in the Netherlands and the rest are widely distributed elsewhere in Europe: 5 in Germany, 1 in France, 1 in England, and 1 in Sweden.

It was first intended as a low frequency telescope to operate below 100 MHz. Early in the design phase, it was decided that it should also search for the Epoch of Reionization. This inspired the extension of the sensitivity band up to 240 MHz. To significantly advance the science below 100 MHz the new instrument needed to be at least 10 times larger than existing facilities. The need to calibrate the effects of the ionosphere was a significant driver for this. Ionospheric effects are discussed further in subsection 3.2.2. Improved ionospheric calibration improves the dynamic range of the resulting measurement. The design was extended to higher frequencies with minimal added complexity and cost in order to search for a signal from the Epoch of Reionization. Cost is minimized by using very simple, stationary antennas.

3.1.2 Target Science

The primary goals which have driven the design of the instrument are low frequency surveying and search for the EoR signal. There are also many other scientific investigations for which LOFAR will be extremely useful.

Early in the expansion of the universe, as it cooled, nuclei captured electrons and became neutral. This is the event we see in the cosmic microwave background (CMB) as this was the first time the Universe was optically transparent. Later when the first stars turned on, ultraviolet radiation reionized the intergalactic medium. This is what is referred to as the Epoch of Reionization. LOFAR will use the 21 cm HI emission line to probe this. The intensity of this radiation dropped in proportion to the neutral fraction of hydrogen, so an increase in intensity with redshift should be seen. The ionized areas expanded spherically about the first light sources, so this

will not be an isotropic signal. LOFAR has been designed to detect the EoR signal in the redshift range z=6-15.[10]

Deep extragalactic surveys will be performed. It is expected that the deepest of these surveys will reach a noise floor of 6μJy at 150 MHz with 5" resolution. Most sources found in these new surveys are expected to be star forming galaxies rather than Active Galactic Nuclei, which most known extragalactic radio sources are. LOFAR should be able to detect star formation rates of order 10 $M_\odot$/yr at redshift z=2. It is anticipated that order 10^{7-8} star forming galaxies should be detected [11].

LOFAR will also be capable of detecting transient sources such as pulsars, supernovae, accreting supermassive black holes, and flare stars. There may also be unknown sources of pulsed radio emission which we have yet to identify. High and Ultra High Energy Cosmic Rays (HECRs and UHECRs respectively) emit a very intense, very brief radio pulse when they enter our atmosphere. Future upgrades could enable detection of HECRs and UHECRs by this mechanism [12].

The sun and space weather are also targets of LOFAR investigation. Flares and coronal mass ejections are important to understand to be able to predict space weather and its impact on electrical systems on Earth as well as for astronaut safety on possible future long-range missions [12].

Magnetic fields may be measured and mapped by the synchrotron radiation they produce and the Faraday rotation they induce. Both of these effects are most significant at low frequencies, so LOFAR will be well suited to measure them. For a brief explanation of Faraday rotation please see footnote 1 on page 46.

3.1.3 Layout, Antennas, and Receivers

There are 48 LOFAR substations. Of these, 22 are designated 'core' stations. These are clustered in a small area 2 km in diameter in Drenthe, The Netherlands. Another 18 stations are also located elsewhere in The Netherlands. These are called 'remote' stations. They are spread up to ~ 100 km from the core. The remaining 8 stations are the 'European' stations. They are at distances of up to 700 km from the core. There are 5 in Germany and 1 each in France, United Kingdom, and Sweden [12].

LOFAR sensitivity reaches from 15 MHz to 240 MHz. Only 48 MHz of this is available at any one time [12]. This is discussed further in subsection 3.1.4. The FM band from 87 MHz to 108 MHz is excluded. Due to this exclusion and since the desired sensitivity band is very broad, it was natural to use two different antenna designs, one optimized for high frequency and the other for low.

The Low Band Antennas (LBA) were optimized for 30 MHz to 80 MHz. They are simple dipoles arranged semi-randomly over a diameter of about 81 m with an exponential decrease in density with radius. Each station has 96 LBA. Theoretical and experimental work indicated that it is not necessary for all stations to have the same beamshape. The core, which is of particular importance for EoR science and ionospheric calibration, has been optimized differently than the remote stations. For core stations, the receivers have two low frequency inputs, each connected to 48 of the antennas. The inputs differ only in the bandpass characteristics (one extends to lower frequencies). The two sets of antennas are arranged slightly differently

Figure 3–1: Dutch LOFAR substation. Note the two clusters of High Band Antennas in the foreground and background, and the quasi-randomly arranged Low Band Antennas in between. There are 24 High Band Antennas in Each cluster and 96 Low Band Antennas in the random dispersion [13].

because of their different target frequencies. In the remote European stations, all 96 are connected to a single low band input.

The High Band Antenna (HBA) was designed for 120 MHz to 240 MHz. It consists of a small array of 16 dipoles arranged in a tight-packed 4 by 4 grid. An analog beamformer combines the signal from all the dipoles, thus forming them into a single 'antenna'. This increases the effective area of the HBA for minimal cost. They have an effective area similar to the LBA. There are 48 in each of the Dutch substations and 96 in each of the non-Dutch ones.

Nyquist's theorem states that all the information in a signal can be perfectly recovered if the sampling frequency (f_s) is greater than twice the largest frequency component of the signal (f_{max})[14]. If this condition is not met, it will not be possible to distinguish certain frequencies. For example, two frequencies $f_1 < f_s/2$ and $f_2 = f_s - f_1$ cannot be separated. This effect is called aliasing. After digitization no distinction can be made between any frequencies of the form $f_n = N \times f_s \pm f_1$ for any integer N[14]. Each frequency range $(K-1) \times f_s < f < K \times f_s$ is called the Kth Nyquist zone. The first three zones are illustrated in figure 3–2 on the following page. If a signal's band is entirely contained in one of these zones, and it is known which zone this is, then the signal can be perfectly reconstructed.

LOFAR's digitization scheme makes extensive use of aliasing. In different observing regimes the first three Nyquist zones are used. The ADC can be clocked at either 160 MHz or 200 MHz. Coupled with the bandpasses of the different inputs, this allows 4 different observing bands to be digitized using the same ADC. The

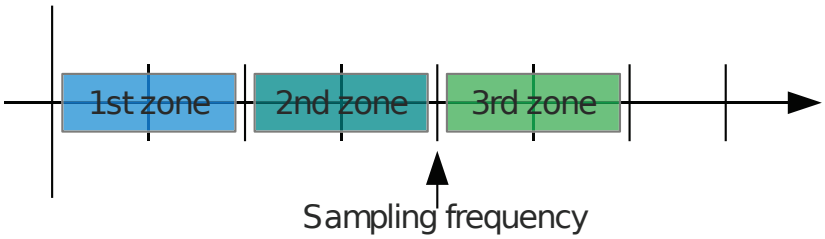

Figure 3–2: Nyquist zones. In most applications the first zone is used and the others are filtered out using low-pass filters. The use of band-pass filters allows the measurement of higher frequencies instead.

bit depth of 12 bits is determined by the extremely bad RFI environment in northern Europe. Each antenna is digitized and beamforming is done digitally. When sampling at 200 MHz this means that 230 Gb/s are produced by each station.

3.1.4 Digital Processing

After digitization, a polyphase filter bank divides the signal into 512 frequency bins of equal width. This is either 156 kHz or 195 kHz depending on choice of digitizer frequency [2]. Up to a total of 48 MHz of these may be selected for further processing [12]. These frequency channels need not be adjacent, and the same frequency channels need not be selected from all stations. A single station cannot use the LBA and HBA simultaneously, but different stations can use different antennas. In spite of filtering out a substantial amount of information, algorithms exist to nearly recreate timestreams for pulsar studies. Calibration is performed using all 512 frequency channels, one at a time. The calibration process for each one

requires ~1 second, so the calibration parameters at all frequencies are updated approximately every 8.5 minutes. A more extensive treatment of LOFAR's calibration algorithm is given in section 3.2.

At this stage phase shift beamforming is employed to 'point' the station. The width of the frequency channelization is sufficiently small to minimize the distortion created by this technique. Each frequency channel is treated separately. This allows beams at different frequencies to be pointed in different directions. The number of beams produced is limited by processing power. A schematic diagram of a European substation is shown in figure 3–3.

At the central computing node, the data is further subdivided into narrower frequency channels. It is channelized into 2^n frequency channels, where n is an integer from 0 to 12. Channelization limits the time resolution in each frequency channel; the time resolution in a frequency channel is inversely proportional to the bandwidth of it: $\triangle t \propto \frac{1}{\triangle \nu}$. For example, if $n = 8$ is selected, which is a typical value, the bandwidth of each frequency bin will be $\Delta\nu = 763$ Hz and the time resolution will be $\Delta t = 1.3$ ms. This reduction in timing resolution can be a problem for pulsar studies, but algorithms exist which will be able to reproduce timestreams for pulsar observations.

Further calibration is performed at the central node based on a global sky model. The difference between the model and measured sky is used to update the calibration parameters and the sky model. Initially known catalogues will be used to generate the model sky. LOFAR's sensitivity and resolution are greatly superior

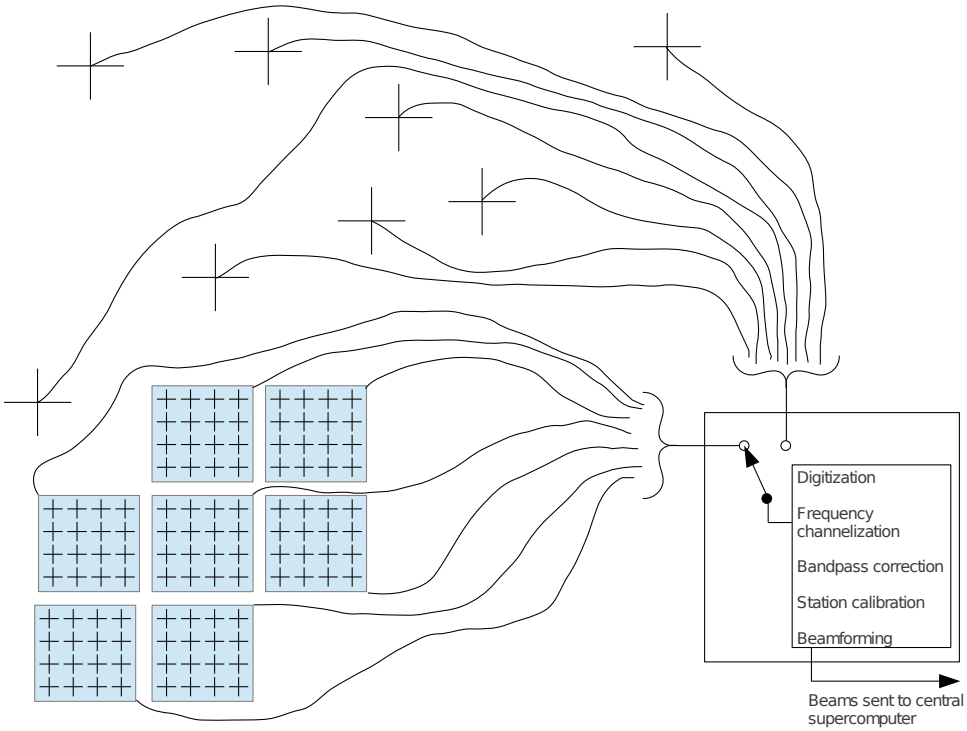

Figure 3–3: European substation schematic. 96 High Band Antennas (closely set) or 96 Low Band Antennas (scattered) are connected to the digital processing circuitry. Only a few of the 96 are shown for clarity. At the substation level the signal is digitized and frequency filtered into 512 channels. Gain corrections are applied to compensate for imperfections in the bandpass filter and the amplifier gains and phases are calibrated. The substations are small, so phase shift beamforming is employed rather than true delays.

to any low frequency instrument before it, so it will discover a vast number of new sources. These will also be incorporated into the sky model.[15]

The raw data rate is 25 PB/day. After processing and time averaging, this is reduced to 250 TB/day. This semi-raw data is stored for up to a week to permit post-processing.

3.1.5 Imaging

Most extant interferometers have a narrow primary beam. This allows the assumption to be made that the array geometry appears constant over the whole Field of View (FoV). In the case of LOFAR and certain other large arrays, this assumption breaks down. This means that the effects of the 'W' (line of sight) component of the array geometry cannot be accounted for by simple phase adjustments.

There are two ways to deal with this difficulty. One is to form the image in facets which are designed to be small enough that the constant projection assumption holds. The other is to apply a w-dependent convolution to the U-V data during griding. This is designed such that a two dimensional Fourier transform will then produce an undistorted image [11]. LOFAR uses the first technique [15].

If the first of these techniques is used, then the Direction Dependent Effects (DDE) must be taken into account separately for each facet. The time and frequency dependence of the corrections also complicates the situation. However, it is possible to solve the DDE in a similar manner to the w-dependent convolution, and indeed to combine the two into a single correction.

3.2 LOFAR Calibration

3.2.1 Introduction

LOFAR calibration and data reduction present several interesting challenges. These include the effects of the ionosphere, which are very pronounced at low frequencies, an extremely bad RFI environment, and substations with significantly different environments from one another. In this section, the calibration techniques used to address these problems will be discussed.

3.2.2 Calibration Issues

LOFAR faces several significant challenges in calibration. These include the ionosphere, beam variations, and bright foreground sources.

The beam of LOFAR has many direction dependent effects. Since it is not steered mechanically but rather is steered electronically by digital beamforming, the beam itself is a function of pointing direction. If a source is tracked, the beam will therefore vary continuously over the course of the observation. The beam is also frequency dependent, but this is generic for all radio telescopes. The interactions between the antennas are very complex, but it may be possible to develop theoretical beam models with 1% accuracy [11].

The ionosphere is a partially ionized zone in the upper atmosphere. The thickness, height, and electron density are all widely variable. There is also significant wind and turbulence in this layer. A phase delay is induced based on the local electron density. This causes distortions in an incoming plane wave. Unlike most telescopes, LOFAR's primary beam is much larger than these distortions, so the

correction for them must be direction dependent. These effects are most significant under 300 MHz [11]. They scale as λ^2.

The low frequency sky is dominated by a few very bright sources. These are Cygnus A, Cassiopeia A, Herculus A, Taurus A, Virgo A, and Sol. LOFAR's beams are wide (~ 10-20 degrees) and have significant sidelobes, so for most pointings one or more of these contributes a significant amount of power. If they are not removed, even from the sidelobes, their power dominates the measurement noise. For mapping purposes, the dynamic range desired in the images is 10,000. In order to achieve this, the calibration in the direction of these sources must be very precise [11].

3.2.3 Measurement Equation

A radio signal can be represented as a 2×1 vector of polarization components. The complex effects of the various problems outlined above may be represented quite simply. Consider a single visibility. Each effect serves to modify the voltage signal which ultimately reaches the correlator. These changes in amplitude and phase can be modelled as a 2×2 matrix (a Jones matrix) operating on an input 'real' vector and giving a 'measured' vector. This can include effects such as polarization leakage, ionospheric phase shifts, transmission losses, and many other effects. The measurement equation illustrates this [11]:

$$V_{pq} = G_p \left(\sum_{i=1}^{N} K_{pi} B_{pi} I_{pi} R_{pi} F_i F_i^{\dagger} R_{pi}^{\dagger} I_{qi}^{\dagger} B_{qi}^{\dagger} K_{qi}^{\dagger} \right) G_q^{\dagger} \tag{3.1}$$

Here V is the visibility between antennas p and q, and G is a Jones matrix which includes only direction independent effects such as gains and polarization leakage.

The sum is over all directions i, and $X_{p(q)i}$ represents various effects of antenna p (or q) in direction i. F_i is the true electric field on the sky. The other $X_{p(q)i}$ represent effects of Faraday rotation ($R_{p(q)i}$), the ionosphere ($I_{p(q)i}$), station beams ($B_{p(q)i}$), and others ($K_{p(q)i}$). This equation is general for any interferometer.

Calibration, in light of the measurement equation, is the attempt to measure the parameters R_{pi}, I_{pi}, B_{pi}, G_{pi}, and K_{pi} so that F_i can be determined accurately. Some effects such as polarization leakage are unlikely to change significantly. In this formalism all distorting effects take a similar form. It is therefore of minimal importance to distinguish between them when calibrating, except to disentangle effects which vary quickly, such as the ionosphere, from those that do so slowly or not at all, such as those due to array geometry.

LOFAR has two levels of hierarchy which differ by many orders of magnitude: the substations, composed of 96 antennas; and the whole system, composed of 48 substations. These differ in scale by a factor of ten thousand. Calibration is carried out at both of these levels, though more care is taken at the system level due to the much greater sensitivity and resolution afforded by the large collecting area and long baselines.

3.2.4 Substation Calibration

Each substation filters the signal to 512 frequency channels. These are calibrated one at a time in about a second, so the whole frequency band is calibrated every 8.5 minutes [2]. RFI is flagged according to its spectral properties and those channels are removed, and the bandpass is corrected for. If the complex voltage

signal of antenna i is v_i then the correlation matrix is $c_{ij} = v_i v_j$. The full correlation matrix of all 96 antennas is used for gain and phase calibration [15]. Each antenna sees the same sky, so the voltages are a direct measure of the gain of each antenna. These are easily determined up to a constant, which can be set by a simple sky model. The phases may similarly be solved for up to an unimportant additive constant by comparing the off-diagonal terms in the correlation matrix. Each off-diagonal term c_{ij}, $i \neq j$ contains the phase difference between antennas i and j. There are 4,560 independent phase difference measurements in the correlation matrix and only 96 phases to determine, so the problem is massively overconstrained. Only 96 of the off-diagonal terms are necessary to measure the phases of all antennas. The power from the bright sources mentioned in subsection 3.2.2 is subtracted at the substation level [2].

Beamforming is performed at the substation level. This is done on each frequency channel separately. They are narrow enough that phase shifts may be used to steer the array rather than geometric time delays. So long as the total bandwidth of 48 MHz is not exceeded, the observation strategy may choose to point the beams in several different directions. several beams may use the same frequency channels, but then the total bandwidth will be reduced. For example, one might choose to form 1 beam from 30 MHz to 78 MHz, or 8 beams from 30 MHz to 36 MHz, or the same 8 beams may have different frequency ranges from one another. The total number of beams is limited by computing power at the substation level. The initial design allows for eight simultaneous beam pointings. The beams are formed by a

complex weighted sum of the voltages. This weighting points the beam in the desired direction by shifting the phase of each antenna appropriately and compensates for the gain and phase differences measured during substation calibration. These formed beams are sent to a central computing node in Groningen, The Netherlands where the Major Cycle of calibration takes place [2].

3.2.5 System Calibration

The system level calibration is based on a global sky model. It will initially be based on presently available catalogues such as the National radio astronomy observatry (NRAO) Very large array (VLA) Sky Survey (NVSS) [15]. A local sky model is built based on specific observational parameters such as observing band and field of view[15]. Six extremely bright sources have already been removed [2]. The data is projected into the u-v plane. The brightest sources remaining in the field and sidelobes are used to make a fit of their parameters (shape, intensity, polarization), instrumental gains, and ionospheric effects [15].

The local sky model and this first crude estimate of the direction dependent calibration parameters are used to predict what the u-v data are expected to be. The difference between this and the actual u-v data will leave some residual. This residual represents errors both in the sky model and the calibration parameters [15].

A correction is made based on current parameters and the residual is converted to image space. This is done in many small facets so that the projection of the array may be taken as constant over the field of view of each facet. Some small distortion remains, which manifests itself as a slightly position dependent point

spread function. Residual power from known sources is used to update the local sky model. New sources will also be found and included in the model [15].

This cycle is iterative. As the solution for the local sky model becomes more robust it is used to refine the global model [15]. The fit of the DDE parameters is also refined at each step. The global sky model will become considerably more complex, as LOFAR is expected to find many millions of new sources [11].

3.2.6 Summary

A telescope's sensitivity is a function of frequency and direction. LOFAR's beams vary as a function of time and have significant sidelobes, which means sources outside the field of view may not be ignored [15]. The measurement equation expresses the problem of calibration as solving for a direction dependent Jones matrix.

Substations are internally calibrated using correlations between its dipoles. RFI is also flagged and removed here. A number of beams are formed using a weighted sum of the power. The overall calibration is accomplished iteratively by predicting expected visibility data, making a fit for calibration parameters based on the data, and updating the models based on the difference between the predicted and measured data.

CHAPTER 4
Murchison Widefield Array

The Murchison Widefield Array (MWA) has been recently built in a radio quiet zone in Western Australia. Its target science includes the Epoch of Reionization (EoR) and the heliosphere. The design of the instrument is described in section 4.1. The MWA is a low frequency instrument sensitive from 80 MHz to 300 MHz. This presents interesting challenges and opportunities, which are discussed. Calibration is performed in real time to enable the data to be stored in the form of integrated images. An innovative approach is used whereby bright point sources either in the main beam or sidelobes are removed from the data set entirely before weaker ones are considered at all. This process is described extensively in section 4.2.

4.1 Introduction to the Instrument

In this section the overall design of the Murchison Widefield Array is described. Unless otherwise noted, all material in this section is derived from "The Murchison Widefield Array: Design Overview" by Colin J. Lonsdale et al [3]. The Murchison Widefield Array is a self-calibrated low frequency dipole based aperture synthesis array in Western Australia.

MWA uses a combination of analog and digital signal processing. The most basic receiving unit is a pair of crossed dipoles. These are arranged in 4x4 close-packed grids (tiles), which are summed in analog. There are 512 of these tiles

spread across a 3km diameter area. The analog summation hardware has various electronically controlled settings which permit pointing of the array across the sky. The array is sensitive from 80-300 MHz, though only 32 MHz of bandwidth may be recorded at a time. At 150 MHz MWA has a primary full width at half maximum (FWHM) beam width of 25^0 and a resolution of 4.5'.

In the 80-300 MHz band there are many sources in the southern sky. There will be a background of many radio bright galaxies within each 4.5' synthesized beam [16]. This inability to distinguish detectable sources results in what is called confusion limited images.

4.1.1 Background

Early interferometry recognized the value of a large range of baselines among the elements. Earth rotation synthesis has often been used to provide this. Only the Giant Meterwave Radio Telescope (GMRT) and the Very Large Array (VLA) currently have sufficient u-v coverage to allow scientifically useful "snapshot" imaging, which refers to forming an image in a short time such that the Earth's rotation is neither significant nor necessary to cover sufficient baselines.

Recent efforts have focused on "large N" arrays which employ hundreds or even thousands of receiving elements. As with single element radio telescopes, the field of view is larger for smaller receiving elements, which provides for faster surveying for a given number of elements. Because of the relationship between collecting area and field of view, Signal to Noise Ratio (SNR) is independent of receiver size for extended sources. A large number of measurements reduces noise on the final result. Therefore, for a given total collecting area, a large number of

small receivers will produce more mapping information than a small number of large receivers. An array of antennas at half-wavelength spacing (Nyquist sampling the electric field of the incoming signal) over the entire collecting area would be ideal, but this is impractical due to the enormous data rate such a system would produce; The size of the data stored on disk (D) is proportional to the bandwidth ($\Delta\nu$) and to frequency (ν) squared: $D \propto \Delta\nu \times \nu^2$. The MWA design attempts to maximize the information content of the data, which optimally constrains both the calibration and the resultant skymap. Measurements of point sources do not gain in sensitivity from large N arrays, except insofar as it may be possible to track multiple sources simultaneously.

4.1.2 Science Objectives

The Epoch of Reionization is the period in the early universe when the first luminous sources began to shine and ionize the intergalactic hydrogen. The MWA will observe this period using the 21 cm hydrogen emission line, which should have decreased in intensity as the gas was ionized. The MWA is designed to detect this signal in the redshift range $6 < z < 20$. This range corresponds to wavelength 1.47 m $< \lambda <$ 4.41 m or frequency 68 MHz $< \nu <$ 204 MHz. The MWA's large field of view increases sensitivity as compared with many other instruments, but even so the collecting area is too small to enable imaging of EoR. Rather, a statistical approach will be used.

Snapshot images of the Sun will be used to study solar radio bursts. Also, 32 independently steerable beams will be formed which will measure interplanetary scintillation. Scintillations are variations in apparent source properties that are due

to variations in some substance through which the radiation passes. Stars twinkling due to turbulence in the atmosphere is an example of this. Here 'interplanetary' refers to the fact that the effect observed is within the solar system, no that the planets themselves are scintillating. Interplanetary scintillation uses the Faraday rotation of background sources to determine the magnetic field present in coronal mass ejections. The substantial effect of Faraday rotation at low frequency is advantageous for these studies. These observations may enable superior understanding and prediction of space weather.

The good Point Spread Function (PSF) and large Field of View (FoV) of MWA are excellent for studying transient sources. Its frequency and temporal resolution are also excellent. MWA will be able to measure variable and transient sources on time scales from tens of nanoseconds to years. It has never before been practical to search such periods on a large scale before. This will therefore open up a new area of astronomy. Multiple back-end systems will conduct both blind and targeted searches.

4.1.3 Low Frequency Challenges

It is physically impractical to build highly directional low frequency antennas, so usually there is significant beam power over the whole visible hemisphere. This has proved difficult in terms of calibration. The ionosphere is also troublesome for an array with such fine resolution (4.5') and low frequency. It can cause refraction and scintillation, and the opacity and induced Faraday rotation are also

variable.[1] These effects generally scale as λ^2. The index of refraction of the ionosphere is dependent on the Free Electron Density (FED) therein. As this density varies, sources appear to move in the sky. Movement of the Earth also causes the signal path through the ionosphere to change in length and angle. This effect is similar to the appearance of a 'broken' spoon in a glass of water. Local variances of the FED are responsible for scintillation [17].

Radio Frequency Interference (RFI) is often a problem, whether from radio broadcasts, airplanes and satellites, or from digital circuitry within the telescope itself. The MWA occupies an extremely radio-quiet area in the outback of Western Australia. Observations indicate a spectral occupancy of ~10^{-4} in 15 second

[1] Faraday rotation is due to the interaction of an external magnetic field with a plasma. The electrons in the field will all gyrate in the same sense with respect to the magnetic field, which results in a different propagation speed for left and right handed polarizations. Equivalently, the index of refraction is polarization dependent. The resultant phase delay between the two circular components has the effect of rotating the polarization of a linearly polarized wave. The rotation angle is given by $\beta = RM\lambda^2$ where β is the rotation angle in radians and RM is the Rotation Measure, given by

$$RM = \frac{e^3}{2\pi m^2 c^4} \int_0^d n_e B_{\parallel} ds \; .$$

The integral is a path integral along the line of sight, n_e is the free electron density, $B_{\parallel}$ is the magnetic field along the line of sight, e and m are the charge and mass of the electron and c is the speed of light. Note that for an observation with a wide field of view, the integral will vary considerably across the observed area [18]. The free electron density is also time variable, especially according to a day/night cycle [17].

integrations at a resolution of 1 kHz and less than 10^{-3} in wider 4 MHz bands integrated for 30-60 minutes. This is negligible contamination. The solution to the difficulties raised by calibration and the ionosphere will be discussed extensively in a later section.

4.1.4 Hardware

The MWA has several major layers of subsystems. In order of sky signal propagation, they are the antennas, the tiles with their analog beamformers, the receiver nodes, the correlator, and the Real Time Computer (RTC) which performs the calibration and forms the images to be saved. These systems must be capable of producing high quality images on a time scale of seconds in order to effectively remove the ionosphere through calibration. These subsystem layers are illustrated diagrammatically in figure 4–1 on the next page.

The antennas selected are dual polarization dipoles with a vertically oriented bowtie profile. This shape has been chosen against a simple dipole for two principal reasons: the antenna pattern is broader, and it improves the impedance matching across the band between the antenna and the Low Noise Amplifier (LNA). This amplifier is located right at the feed point for the antenna making it an active dipole. In creating the bowtie shape only an outline with horizontal crosspiece is used rather than a solid sheet of metal to save cost and weight. The sky noise dominates over system temperature by a factor of several over most of the bandwidth and in most pointing directions. After amplification the balanced signal is converted to an unbalanced one and transmitted to the tile beamformer over 7 m of coaxial cable. The DC power for the LNA is transmitted over the same cable.

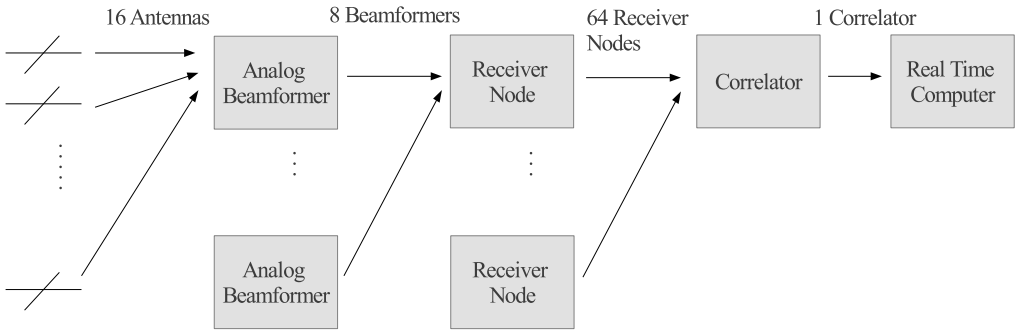

Figure 4–1: Diagram showing the hierarchy of MWA subsystems. The 16 antennas of each polarization on each tile are added, possibly using some delay lines to point the beam, to form a single beam for each tile. Thus each tile functions like a single dual polarized antenna. The two polarizations from 8 antenna tiles are digitized and channelized in the receiver node. The correlator correlates all signals from all 64 receiver nodes. Analysis by the Real Time Computer (RTC) determines correlation parameters, and the data stream is converted to images and integrated.

These dual polarization dipoles are arranged in 4×4 grids over a 5 m $\times$ 5 m steel mesh ground plane. They are spaced 1.1 m on centre, which is an idealized $\frac{\lambda}{2}$ spacing at 136 MHz. An antenna tile is shown in figure 4–2 on page 50. At all frequencies below this, the incoming radiation is oversampled. The mesh groundplane is used as a reference grid for each tile, allowing them to be constructed identically to a 1 cm tolerance. Beamforming is accomplished by mechanical means. Each antenna has a system of five switchable delay lines which may be independently set, for a total of 32 possible delays, including zero delay, for each polarization of each antenna. The output from all the delay lines of a polarization are added passively. The resultant effective tile beam is 25° wide at 150 MHz [16]. Walsh functions are a

certain set of functions which take the value ± 1 at all points. They form an orthogonal basis which spans the set of all real-valued functions [19]. Multiplication by Walsh functions provides a means to distinguish the 16 signals which arrive at each receiver node. This is mathematically similar to encoding different radio stations on different carrier frequencies. The encoded signals are amplified again and sent to a receiver node over RG6 cables. Power and waveforms are sent to the beamformers over the same cables. Since the only signal injected into the beamformer is the Walsh functions around 100 kHz, far outside MWA's sensitivity band, there is no need for elaborate RFI shielding measures.

Each receiver node receives 16 signals in total: 2 polarizations from 8 antenna tiles. The final analog signal shaping occurs here. The signal passes through an antialiasing filter and a highpass filter, signal level adjustments are made, and cable losses are equalized if necessary. The input signals are then digitized by ATMEL 8bit ADCs running at 655 MHz. A polyphase filterbank instantiated on a Field Programmable Gate Array (FPGA) divides the signal into 256 channels of 1.28 MHz width. To reduce the data rate to a manageable level, only 24 of these channels are selected for further processing. The rest are discarded. MWA thus has an instantaneous bandwidth of 30.72 MHz. The selected frequency channels are sent to the correlator over three optical fibres. Each carries one third of the 24 channels. The data are 5+5 bit complex samples at this stage. Monitoring and control and timing signals are sent to the receiver nodes over separate fibres. Each receiver node and its dependent antenna tiles are controlled by a single master board. This stage of

Figure 4–2: Photograph of an MWA antenna tile. 16 dual polarization 'bowtie' dipoles are arranged in a 4×4 pattern. The wire mesh serves as a ground plane and also provides a reference grid to ensure each tile is laid out identically [20].

processing could produce significant RFI, but these nodes are well shielded. The aggregate data rate from all the receiver nodes is over 300 Gbit/s.

At the correlator, another 16 FPGA boards further channelize the signal to just over 3000 10 kHz channels. Graphics Processing Units (GPUs), computer processing chips designed to deal with the vector and matrix based math demanded by graphics rendering, are well suited to performing the actual correlation. They are considerably more time efficient for the calculations needed here [21]. Extensive time and frequency domain multiplexing is employed at this stage. There are 512 antenna tiles with 2 polarizations, so there are $0.5 \times (2 \times 512)^2 = 524,288$ distinct signal pairs to be correlated. The factor of 0.5 avoids counting each pair twice. Each correlation is integrated for 0.5 seconds. It is common in large arrays to employ geometric delays to ensure that the same wavefront at each antenna that is correlated. However, because of the relatively small size and long wavelength at MWA, the correlation length is sufficient to neglect this altogether, greatly simplifying the correlator. The phase centre does not track a target, but remains constant for each array pointing. This results in the simplest possible correlator in design, hardware, firmware, and control. The complexity is moved downstream to later software analysis. The cost of this is a very high output data rate.

4.1.5 Antenna Distribution

There are several factors which influence the distribution of antenna tiles. On a broad scale, the size must be sufficient to permit calibration and the science goals. There are many prohibited locations due to local vegetation and topography. Within these constraints, the driving impetus is image quality. In an array with so many

elements, it is impractical to find the truly optimal solution. Instead, a configuration that produces sufficient quality images is deemed acceptable.

The result is a pseudorandom distribution with an extremely dense core. Outside the inner 50 m the density of antennas falls off as $1/r^2$. The central circle of diameter 1.5 km contains 496 tiles. The annulus surrounding this from $1.5\,\text{km} < D < 3\,\text{km}$ contains a further 16 tiles widely distributed. This results in a u-v plane that is nearly Nyquist sampled at low wavenumber, and quite well sampled elsewhere.

4.1.6 Calibration and Imaging

The data rate from the correlator is19 GB/s [16]. In MWA's initial configuration frequency averaging will reduce this by a factor of 4 at this stage. This is still too great a rate to be able to practically store the visibilities. MWA will therefore store data in the form of images, integrated for 8 seconds. To form these images it is necessary to calibrate in real time. This section will outline the problem briefly. A more complete treatment of the problem and its solution may be found in section 4.2 on the next page.

The ionosphere is a highly ionized region of the atmosphere extending from 100 km to 1000 km. The ionized gas interacts with low frequency electromagnetic radiation inducing refraction and phase delays which scale as λ^2. The full field of view of MWA is 25 degrees wide, so a huge volume of the ionosphere is in view at any time. This volume is much larger than the characteristic size of turbulent regions and other irregularities in the ionosphere, so the effects are direction dependent. The ionospheric irregularities are, however, larger than the longest baseline of 3

km. Therefore, it may be taken as approximately true that every antenna tile has the same view through the ionosphere in any given direction. This significantly simplifies the calibration problem. The effect of the ionosphere is to produce a shift in the apparent direction of a source, but does not magnify or diminish its brightness. The effect can therefore be removed by a linear phase ramp.

The gain of each tile will vary with time. The central limit theorem suggests that the average phase of all the tiles will be stable. The magnitude will be similarly stable up to an (easily measured) additive constant. Therefore one might calibrate each tile by comparing it with the average of the other 511. The MWA calibration scheme is conceptually similar to this, though its mathematical form appears quite different. The calibration is done on the visibilities rather than the more intuitive voltage stream, which saves computation.

4.2 Calibration of the Murchison Widefield Array

4.2.1 Introduction

Cross-correlations between tiles (visibilities) measure points in Fourier space (u,v,w). For most telescopes phase adjustments can be used to effectively eliminate w, so that the problem reduces to a two dimensional form. However, due to the wide instantaneous field of view, MWA cannot do this. Nonetheless, other techniques have been developed which permit a two dimensional form of the calculation. Effects due to the ionosphere and antenna beam differences also mean MWA cannot calibrate by iterating between image space and visibilities [16].

The MWA tiles have different direction dependent responses which vary with time, in part due to the ionosphere. Bright calibration sources distributed across the

whole field of view, including sidelobes, are used to measure phase shifts induced by the ionosphere and the complex gains of the antenna tiles.

Since the data rate from the MWA correlator was deemed too large to store,images of the sky will be stored instead. This greatly lowers the data rate, but requires that instrument calibration and image formation be done in real time.

Several factors render this process difficult. The analog beams will be updated every 5-10 minutes for source tracking. The ionosphere varies on time scales of ~1 minute. Furthermore, the gain and polarization response of the antenna tiles are strong functions of direction. Galactic synchrotron radiation is a particular issue owing to its ubiquity and intensity, both in the main beams and the sidelobes. The algorithm presented hereafter has been tested in an unpolarized form with simulations[16] and, in a slightly simplified form, using the 32 element prototype [22].

4.2.2 Visibilities

In the following analysis, lower case will indicate a vector, upper case a matrix, and † a Hermitian transpose. The correlator will have M frequency channels with bandwidth Δf which will be integrated for time Δt. The normal situation for MWA is $M = 768$, $\Delta\text{f} = 40$ kHz, $\Delta\text{t} = 8$ s. The integration time of eight seconds has been chosen to oversample the variations of the ionosphere [16].

Most of the power in the visibilities is due to sidelobes, backgrounds, our own galaxy, instrumental noise, and a number (N_c) of bright point sources which will be used for calibration. Let us define the following variables. Each entry in a vector corresponds to a polarization component, either linear or circular.

$p_{c,f}$ = true voltage signal from calibrator 'c' in the frequency channel centred at 'f,' in sky coordinates.

$J_{j,c,f}$ = transfer function from sky polarization coordinates to receiver polarization coordinates of tile 'j' for calibrator 'c' in the frequency channel centred at 'f'.

$r_{j,c,f}$ = response of antenna tile 'j' to calibrator 'c' in the frequency channel centred at 'f,' in ground coordinates.

$P_{c,f}$ = autocorrelation of the signal from calibrator 'c'in the frequency channel centred at 'f'.

$R_{jk,c,f}$ = visibility of tiles 'j' and 'k' in response to calibrator 'c' in the frequency channel centred at 'f'.

The following relations hold for any time scale for which J is approximately constant.

$$r_{j,c,f} = J_{j,c,f} p_{c,f}$$

$$P_{c,f} = < p_{c,f} p_{c,f}^{\dagger} >$$

$$R_{jk,c,f} = < r_{j,c,f} r_{k,c,f}^{\dagger} >$$

$$R_{jk,c,f} = < (J_{j,c,f} p_{c,f})(p_{c,f}^{\dagger} J_{k,c,f}^{\dagger}) >$$

$$R_{jk,c,f} = J_{j,c,f} P_{c,f} J_{k,c,f}^{\dagger}$$

The total power in the visibility jk is given by

$$V_{jk,f} = N_{jk,f} + \sum_{c=1}^{N_c} R_{jk,c,f} e^{\iota \varphi_{jk,c,f}} \tag{4.1}$$

where $N_{jk,f}$ is the noise, both thermal and confusion, and the factor $e^{-\iota\varphi}$ accounts for the differing phase relations in the direction of each calibrator.

4.2.3 Calibrator Measurement Loop

Only point sources are used for calibration. The Molongolo Reference Catalogue has 7347 sources in a usable declination range with a flux over 1 Jy at 408 MHz. Spectral information is not available across the entire MWA band, so some bootstrapping is necessary. To minimize communication between compute nodes, the calibration is parallelized over frequency, so that each node calculates the calibration for K consecutive frequency channels.

The calibration procedure proceeds as follows. The prior set of beam models is used to calculate and rank the power contribution of each calibration source. The total power predicted based on the old beam models is subtracted from each visibility. Then the calibration proceeds in a loop over the calibration sources, in order of their ranked power:

1. The power from the first calibrator is returned to the data set. The phase centre is set to the predicted location of the calibrator. In essence, the telescope 'looks' at the calibrator. Each node averages the signal over the frequency channels it is calibrating.
2. The flux in the data set is now primarily due to a single source. In this step only the ionospheric phase shift is calculated, so polarization information is unnecessary and only the intensity is used. There will be a small deviation between the predicted location and observed location of the calibrator. This

shift is dependent on λ^2. The noise contribution from the foregrounds, backgrounds, and hardware is expected to be nearly normally distributed. Each node sends its measured shift to a central CPU which calculates a phase ramp from the whole data set, returning the result to each node.

3. The ionosphere-corrected calibration is now used to determine the gain of each element. This is essentially a measurement of each $J_{j,c,f}$. This is accomplished by a least-squares minimization of the difference between each measured visibility and that predicted based on the known source.

$$J_{j,c,f_0} = \left(\sum_{k,k\neq j}^{N_a} V^{(K)}_{jk,c,f_0} J_{k,c,f_0} P^{\dagger}_{c,f_0} \right) \times \left(\sum_{k,k\neq j}^{N_a} P_{c,f_0} J^{\dagger}_{k,c,f_0} P^{\dagger}_{c,f_0} \right)^{-1} \tag{4.2}$$

   Much of the information in the Jones matrix may be known in advance, such as the conversion from sky to ground polarization coordinates. The above equations may easily be modified such that only the truly unknown portion is fit.

4. The newly determined gain and ionosphere model are used to fully subtract the calibrator from the data set. This is referred to as 'peeling.' Calibration continues with the next most significant calibrator.

Strong sources are removed from the data set before weaker sources are measured. As the strengths of the calibrators span many orders of magnitude, this is extremely valuable. It should be possible to peel down to a noise floor of a few 10^2 mJy, limited by power in sidelobes. If sources at less than 5 sigma above the noise floor of a single image are used it would be possible to peel more deeply, but at the price of accuracy.

4.2.4 Performance

Many array calibration techniques are known. Most scale as N_a^3, where N_a is the number of digitized receiver elements. The most efficient technique known, the LOGarithmic Least Squares algorithm (LOGLS), scales as N_a^2. Specifically, LOGLS has $O(32N_a^2)$ complex multiplications [23]. The technique presented here has $O(36N_a^2)$, so it compares favourably with known efficiency limits. Nonetheless, the computational demands are considerable: 180 megaFLOP per frequency channel per calibrator. In all, several petaFLOP will be required every 8 seconds [16].

Several possibilities are being considered to reduce this. The ionospheric measurements consist of detecting a small position error. Thus, the longest baselines are most significant to the calculation. It may be possible to ignore 99% of the short baselines without compromising accuracy substantially. Furthermore, the tile beams should vary slowly with time, so it may not be necessary to calibrate fully every 8 seconds. Updating the calibration on 20-30 calibrators at a time should be sufficient. If beam models become sufficiently accurate it may be possible to fit for the gains using all the calibrators at once, but it will not be clear until the instrument has been running for some time whether this will be viable.

Calibration errors are a concern during foreground subtraction as they will be averaged into the stored images. The calibration parameters will be stored, so that errors can be reduced in post-processing. If concern arises that peeling errors may mimic EoR then images may be created using the unpeeled visibilities. The MWA site is a radio quiet zone, but some RFI will inevitably be present. Corrupted

frequency bins will be flagged and ignored, with little impact on either calibration or image quality.

4.2.5 Simulations

Simulations of the entire data reduction process from sky signal to images have been performed. This was done using MIT Array Performance Simulator (MAPS). The fully polarized form of MAPS was not fully tested, so the unpolarized form was used. The movement of ionospheric disturbances has also not been included. At the resolution of MWA, an otherwise very realistic simulation is possible. Both diffuse and point-like emission are included in the simulated sky, as well as the ionosphere. The International Reference Ionosphere is used, with a spectrum of turbulence. The gain of each dipole in each tile has some complex noise, which results in different analog beams for each tile. The system temperature was taken to be 200 K.

It was found that the RMS of the gain error improved only negligibly after ~10 peel iterations. This implies that it will not be necessary to use every source for calibration at every step. It was found that the fits were better if preference was given to sources in the sidelobes rather than ranking strictly based on total power expected in the visibilities.

4.2.6 Summary

The MWA correlator receives signals from 512 dual polarization tiles. An iterative approach to calibration which goes back and forth between u-v space and image space would be computationally prohibitive. Instead the expected contribution from each calibrator source is projected into u-v space and the calibration is performed there. For an instrument of this scale the ionosphere induces an apparent

position shift, but does not magnify power. This phase shift is removed first, then the complex gains in each direction are measured. Calibration is performed first on the brightest sources and then their power is removed. This greatly reduces the impact of power from bright sources contaminating measurements of weaker ones through calibration errors [3].

Long term, only image space data is stored. After all the calibration and bright source removal has been performed for each eight seconds of integrated u-v data, it is converted into an image. The images are further integrated to reduce data storage requirements [3].

CHAPTER 5
Conclusion

5.1 Instrument Comparison

Two important modern instruments and their calibration schemes have been discussed. Further insight may be gained by their comparison.

5.1.1 Science

MWA and LOFAR were both designed with EoR in mind. LOFAR was also intended as a low frequency survey instrument. They each have a similar range of secondary science available to them such as the study of transients, the search for new transients or new kinds of transient sources, solar observations, and measurements of magnetic fields through Faraday rotation. In spite of these similarities of target science, MWA and LOFAR are quite different.

5.1.2 Hardware

The scale of these instruments is quite different; LOFAR has a maximum baseline of 700 km, but MWA is only 3 km in diameter. Thus its resolution is also less by a factor of ~100. This has an enormous impact on the ability to discern a point source over sky noise. For a given collecting area, this aspect of the sensitivity varies as the square of the resolution. The sensitivity is also proportional to the collecting area, and LOFAR is much larger than MWA. MWA will not be able to see the many millions of point sources that LOFAR will. However, this effect which

reduces point source sensitivity has no effect on diffuse source sensitivity. The EoR will manifest as a diffuse signal, so this is no obstacle to MWA's primary science.

Both instruments use active dipoles as a receiving element. However LOFAR uses different antennas for different parts of its band. Both these antennas cannot be used simultaneously at a single substation so LOFAR has strict constraints on the subbands chosen for analysis which MWA does not share. LOFAR also blocks out the FM band entirely. The LOFAR High Band Antenna (HBA) is actually a grid of 16 dipoles summed in analog. It is thus similar to an MWA tile, except that the HBA beam is not steerable as the MWA tile beams are.

MWA has 512 antenna tiles. The calibration and data reduction for all of them is conducted at a single central location. LOFAR, on the other hand, has another layer in the hierarchy of subsystems: the substation. These comprise 96 antennas of each band. They are calibrated for gain and phase locally. They are digitally beamformed, forming one or more beams. These signals are then combined in a central location where further calibration takes place.

5.1.3 Calibration

The calibration schemes used by these instruments differ considerably.

LOFAR substations use the correlation matrix from all 96 antennas to calibrate their complex gains. The contributions of the brightest sources in the sky are then subtracted from the visibilities before beamforming. A second stage of calibration using the whole of LOFAR then deals with atmospheric distortions, which may be different for different substations, and more precise gain measurements. This is where direction-dependent effects are first taken into account. This calibration

functions by comparing a local sky model to the visibilities actually measured. Iterative adjustments to the sky model and direction dependent calibration parameters continue until the quality of the match is sufficient. This is conducted one frequency at a time. It takes a total of about 8.5 minutes to calibrate all frequencies.

In contrast, MWA iterates over calibration sources. Calibration is parallelized over frequency. The calibration is calculated entirely on the raw visibility data to avoid repeating the computationally expensive conversion to image space more than is necessary for image integration. The calibrator's power is projected into visibility space. Recent calibration information is used to remove all of the calibration sources from the data before any calculations take place. Then the power from each source is returned one at a time and calibration parameters for ionospheric phase shift and gain fluctuations are determined. The power from the source can then be removed from the visibilities more accurately. In this way, each calibration source is the dominant source of power whenever it is being used, even if it is one of the weaker sources. This calibration loop can be completed in approximately 8 seconds, much faster than LOFAR. This is partially due to the reduced precision necessary because of lower resolution. The lack of iteration between u-v and image space also saves a tremendous amount of computing power/time.

5.2 Conclusion

LOFAR has begun observations of some of the most foreground free viewing directions where the EoR will be sought. The EoR itself has not yet been detected,

but calibration routines seem to be working well and there appears to be no great obstacle remaining other than the need for long observation times [24]. Pulsar studies have already begun with two new discoveries[25].

Solar observations have been carried out with 32 element MWA prototypes. These have revealed a large variety of weak transient events [26]. Observations have also been carried out on sources known well at higher frequencies [27]. Construction continues on a 128 tile version of MWA which is expected to be able to detect the EoR [28].

In this thesis an introduction to radio astronomy and interferometry has been given. New low frequency instruments, requiring new calibration techniques, are being developed, spurred on by increasing computer power. These developments have been studied in the context of the LOw Frequency ARray (LOFAR) and the Murchison Widefield Array (MWA). The techniques pioneered by these instruments may become typical as a new generation of instruments is built.

REFERENCES

[1] EVN collaboration. http://www.evlbi.org/intro/intro.html. Accessed 27 August 2013

[2] van Haarlem, M. P. "LOFAR: The LOw Frequency ARray." Astronomy & Astrophysics 56.A2, (2013):.

[3] Lonsdale, Colin J., et al. "The Murchison Widefield array: design overview." Proceedings of the IEEE 97.8 (2009): 1497-1506.

[4] Ian R. Kenyon, **The Light Fantastic: A Modern Introduction to Classical and Quantum Optics**, Oxford University Press, Oxford, 2008.

[5] Carroll, Bradley W., and Dale A. Ostlie. **An Introduction to Modern Astrophysics** (2nd edition). Addison-Wesley, Boston, 2006.

[6] Gerrald L Pollack and Daniel R Stump, **Electromagnetism**, Addison-Wesley, Boston, 2001.

[7] Orfanidis, Sophocles J. **Electromagnetic Waves and Antennas**. Rutgers University Press, 1999. (electronic version accessed from http://www.ece.rutgers.edu/~orfanidi/ewa/ on 13 September 2013)

[8] Thompson, A. R., J. M. Moran, and G. W. Swenson Jr. **Interferometry and Synthesis in Radio Astronomy** (2nd edition). WILEY-VCH, Weinheim, 2004.

[9] de Vos, Marco, Andre W. Gunst, and Ronald Nijboer. "The LOFAR telescope: System architecture and signal processing." Proceedings of the IEEE 97.8 (2009): 1431-1437.

[10] Zaroubi, Saleem. "The Epoch of Reionization." Astrophysics and Space Science Library. Vol. 396. 2013.

[11] Tasse, Cyril, et al. "LOFAR calibration and wide-field imaging." Comptes Rendus Physique 13.1 (2012): 28-32.

[12] GRIESSMEIER, Jean-Mathias, Philippe ZARKA, and Michel TAGGER. "Radioastronomy with LOFAR." Comptes rendus. Physique 13.1 (2012): 23-27.

[13] LOFAR collaboration. Downloaded from https://www.astron.nl/sites/astron.nl/files/cms/lofar_2010-05-23_1.jpg on 12 September 2013.

[14] Marks, Robert J. II. **Introduction to Shannon Sampling and Interpolation Theory**. Springer-Verlag, New York, 1991.

[15] Nijboer, R. J., and J. E. Noordam. "LOFAR calibration." Astronomical Data Analysis Software and Systems XVI. Vol. 376. 2007.

[16] Mitchell, Daniel A., et al. "Real-time calibration of the Murchison Widefield Array." Selected Topics in Signal Processing, IEEE Journal of 2.5 (2008): 707-717.

[17] Williams, Christopher L., et al. "Low-frequency Imaging of Fields at High Galactic Latitude with the Murchison Widefield Array 32 Element Prototype." The Astrophysical Journal 755 (2012): 47.

[18] Rybicki, George B., and Alan P. Lightman. **Radiative Processes in Astrophysics** (2nd edition). WILEY-VCH, Weinheim, 2004.

[19] Harmuth, Henning F. **Transmission of Information by Orthogonal Functions** (2nd edition). Springer-Verlag, New York, Heidelberg, Berlin, 1972.

[20] MWA collaboration. Downloaded from http://www.mwatelescope.org/plugins/content/sige/plugin_sige/download.php?img=%2Fimages% 2FRotatorImages%2Ffront_001.jpg on 19 June 2013.

[21] Ord, S., et al. "Graphics Processing Units for Data Processing in the Murchison Wide-field Array." Astronomical Data Analysis Software and Systems XVIII. Vol. 411. 2009.

[22] Ord, S. M., et al. "Interferometric imaging with the 32 element Murchison wide-field array." Publications of the Astronomical Society of the Pacific 122.897 (2010): 1353-1366.

[23] Boonstra, A-J., and A-J. Van der Veen. "Gain calibration methods for radio telescope arrays." Signal Processing, IEEE Transactions on 51.1 (2003): 25-38.

[24] Yatawatta, S., et al. "Initial deep LOFAR observations of epoch of reionization windows. I. The north celestial pole." Astronomy and Astrophysics 550 (2013): 136.

[25] Coenen, Thijs. "Pilot pulsar surveys with LOFAR." Proceedings of the International Astronomical Union 8.S291 (2012): 229-232.

[26] Tingay, S. J., et al. "The Murchison Widefield Array: The Square Kilometre Array Precursor at Low Radio Frequencies." Publications of the Astronomical Society of Australia 30 (2013): e007.

[27] McKinley, B., et al. "The giant lobes of Centaurus A observed at 118 MHz with the Murchison Widefield Array." Monthly Notices of the Royal Astronomical Society 436.2 (2013): 1286-1301.

[28] Beardsley, A. P., et al. "The EoR sensitivity of the Murchison Widefield Array." Monthly notices of the Royal Astronomical Society. Letters 429.1 (2013): L5-L9.

Glossary of Acronyms

ADC - Analog to Digital Converter

CCD - Charge-Coupled Device

CPU - Central Processing Unit

DC - Direct Current

DDE - Direction Dependent Effect

DFT - Discrete Fourier Transform

EM - Electromagnetic

EoR - Epoch of Reionization

FED - Free Electron Density

FoV - Field of View

FFT - Fast Fourier Transform

FLOP - FLoating-point OPeration

FPGA - Field Programmable Gate Array

FWHM - Full Width at Half Maximum

GMRT - Giant Meterwave Radio Telescope

GPU - Graphics Processing Unit

HBA - High Band Antenna

HECR - High-Energy Cosmic Ray

LBA - Low Band Antenna

LNA - Low Noise Amplifier

LOFAR - LOw Frequency ARray

LOGLS - LOGarithmic Least Squares (algorithm)

MAPS - MIT Array Performance Simulator

MWA - Murchison Widefield Array

NRAO - National Radio Astronomy Observatory

NVSS - NRAO VLA Sky Survey

PSF - Point Spread Function

RFI - Radio Frequency Interference

RM - Rotation Measure

RMS - Root Mean Square

RTC - Real Time Computer

SNR - Signal to Noise Ratio

UHECR - Ultra-High-Energy Cosmic Ray

VLA - Very Large Array

VLBI - Very Long Baseline Interferometry